arXiv:1610.03774v2 [stat.ML] 27 Oct 2016

# Parallelizing Stochastic Approximation Through Mini-Batching and Tail-Averaging

Prateek Jain*, Sham M. Kakade†, Rahul Kidambi‡, Praneeth Netrapalli§, Aaron Sidford¶

**Abstract**

This work characterizes the benefits of averaging techniques widely used in conjunction with stochastic gradient descent (SGD). In particular, this work sharply analyzes: (1) mini-batching, a method of averaging many samples of the gradient to both reduce the variance of a stochastic gradient estimate and for parallelizing SGD and (2) tail-averaging, a method involving averaging the final few iterates of SGD in order to decrease the variance in SGD's final iterate. This work presents the first tight non-asymptotic generalization error bounds for these schemes for the stochastic approximation problem of least squares regression.

Furthermore, this work establishes a precise problem-dependent extent to which mini-batching can be used to yield provable near-linear parallelization speedups over SGD with batch size one. These results are utilized in providing a highly parallelizable SGD algorithm that obtains the optimal statistical error rate with nearly the same number of serial updates as batch gradient descent, which improves significantly over existing SGD-style methods.

Finally, this work sheds light on some fundamental differences in SGD's behavior when dealing with agnostic noise in the (non-realizable) least squares regression problem. In particular, the work shows that the stepsizes that ensure optimal statistical error rates for the agnostic case must be a function of the noise properties.

The central analysis tools used by this paper are obtained through generalizing the operator view of averaged SGD, introduced by Défossez and Bach [1] followed by developing a novel analysis in bounding these operators to characterize the generalization error. These techniques may be of broader interest in analyzing various computational aspects of stochastic approximation.

## 1 Introduction and problem setup

With the ever increasing size of modern day datasets, practical algorithms for machine learning are increasingly constrained to spend less time and use less memory. This makes it particularly desirable for simple streaming algorithms that generalize well in just a few passes over the dataset. Stochastic gradient descent (SGD) is perhaps the simplest and most well studied algorithm that meets these constraints. The algorithm simply repeatedly samples an instance from the stream of data and updates the current parameter estimate using the gradient of the sampled instance. Nevertheless, SGD has been impressively successful and is the de-facto method of choice for large scale learning problems.

While a powerful machine learning tool, unfortunately SGD in its simplest forms is inherently serial. Over the past years as dataset sizes have grown there have been remarkable developments in processing capabilities, with multi-core/distributed/GPU based computing infrastructure available in relative abundance. The presence of this computing power has triggered the development of parallel/distributed machine learning algorithms [2, 3, 4, 5, 6, 7] that possess the capability to utilize multiple cores/machines. However, despite this exciting work, it is yet unclear how to best parallelize SGD and fully utilize these computing infrastructures.

This paper takes a step towards answering this question, by characterizing the behavior of constant stepsize SGD for the problem of strongly convex stochastic least square regression (LSR) under two averaging techniques widely believed to improve the performance of SGD. In particular, this work considers the natural

*Microsoft Research, India, e-mail: prajain@microsoft.com
†University of Washington, Seattle, e-mail: sham@cs.washington.edu
‡University of Washington, Seattle, e-mail: rkidambi@uw.edu
§Microsoft Research, India, e-mail: praneeth@microsoft.com
¶Stanford University, Palo Alto CA, e-mail: sidford@stanford.edu

parallelization technique of *mini-batching*, where multiple points are processed simultaneously and the current iterate is updated by the average gradient over these points, and combine it with variance reducing technique of *tail-averaging*, where the average of many of the final iterates are returned as SGD's estimate of the solution.

In this work, the arguments about parallelization are structured through the lens of a *work-depth* tradeoff: *work* refers to the total amount of computation required to reach a certain generalization error, and *depth* refers to the amount of serial computation. Depth, defined in this manner, is equal to the runtime of the algorithm on a large multi-core architecture with shared memory, where there is no communication overhead, and has strong implications for parallelizability on other architectures as well.

**The main contributions of this paper** are as follows:

- This work shows that mini-batching yields near-linear parallelization speedups over the standard serial SGD (i.e. with batch size 1), as long as the mini-batch size is smaller than a problem dependent quantity (denoted by $b_{\text{thresh}}$). When batch-sizes increase beyond $b_{\text{thresh}}$, mini-batching is inefficient (owing to the lack of serial updates), thus obtaining only sub-linear speedups over and above mini-batching with a batch size $b_{\text{thresh}}$. A by-product of this analysis sheds light on how the step sizes naturally interpolate from ones used by standard serial SGD (with batch size 1) to ones used by batch gradient descent.

- While the final iterate of SGD decays initial error at a geometric rate but does not obtain optimal statistical error rates, the average iterate [8, 1] decays the initial error at a sublinear rate while preserving statistical optimality. This work rigorously shows that tail-averaging obtains the best of both worlds – decaying initial error at a geometric rate and obtaining near-optimal statistical error rates.

- Combining the above results, this paper provides a mini-batching and tail-averaging version of SGD that is highly parallelizable – the number of serial steps (which is a proxy for the un-parallelizable time) of this algorithm nearly matches that of *offline gradient descent* and is lower than the serial time of all existing streaming LSR algorithms. See Table 1 for comparison. We note that these results are obtained by providing a tight finite-sample analysis of the effects of mini-batching and tail-averaging with large constant learning rate schemes.

- All the results in this paper are established for the *general agnostic noise* case of the streaming LSR problem. This reveals a fundamental difference in the behavior of SGD when dealing with the agnostic case as opposed to the realizable case. In particular, this analysis reveals a surprising insight that the maximal stepsizes that are sufficient to obtain statistically optimal rates are a function of the properties of the agnostic noise. The main result of this analysis is that the maximal step sizes for the agnostic case could be *much lower* than ones that are used in the realizable case: indeed, an instance that yields such a separation between the maximal learning rates for the realizable and agnostic case is provided.

The tool employed in obtaining these results is Algorithm 1 and its generalization guarantees as in Theorem 1, which builds upon and generalizes the analysis of [1] and provides the first tight bound on the finite-sample generalization error for mini-batching, tail-averaged constant step-size SGD. We note that the work of [1] does not establish statistically optimal rates while working with large constant step size schemes; this shortcoming is remedied through coming up with a novel sharp analysis that is used to rigorously establish statistically optimal rates while working with large constant step sizes (this can be viewed as a special case of Theorem 2 by setting batch size $b = 1$ and through averaging iterates from the start by setting $s = 0$).

Applying Theorem 1 to the general agnostic case of LSR, we obtain Theorems 2 and 3, which shows an interesting trade-off between the convergence rate or generalization error (work) of SGD, the amount of parallelization (depth of computation) achieved by increasing the mini-batch size, and the amount of tail-averaging performed. As with previous SGD-style algorithms, [8, 9, 10, 11, 1, 12], this trade-off stems from the two terms of the generalization error of SGD: (i) the *bias* term which represents how quickly the initial conditions are forgotten, and (ii) the *variance* term which represents the dependence of the excess risk on the level of noise (due to streaming data/random sampling/model mismatch) that is present in the problem. We note that existing results [9, 1] dealing with averaged SGD with batch size 1 applied to the realizable case of LSR generalize to this paper's results which establishes (a) the finite-sample effects of mini-batching, tail averaging in the context of stochastic approximation, and (b) the impact of agnostic noise on the learning rates employed by SGD to guarantee statistically optimal error rates.

The results of this paper are obtained through building on and generalizing the operator view of averaged SGD with batch size one, as introduced by [1] to the mini-batched tail-averaged case and through subsequently

| Algorithm | Final error | Runtime/Work | Depth | Streaming | Agnostic |
|---|---|---|---|---|---|
| Batch gradient descent | $\mathcal{O}\left(\frac{\sigma^2 d}{n}\right)$ | $\kappa n d \log \frac{n \cdot \text{init}}{\sigma^2 d}$ | $\kappa \log \frac{n \cdot \text{init}}{\sigma^2 d}$ | $\times$ | ✓ |
| Stochastic Dual Co-ordinate Ascent [13] | $\mathcal{O}\left(\frac{\sigma^2 d}{n}\right)$ | $(n+\frac{R^2}{\lambda_{\min}}d)d \cdot \log \frac{n \cdot \text{init}}{\sigma^2 d}$ | $(n+\frac{R^2}{\lambda_{\min}}d) \cdot \log \frac{n \cdot \text{init}}{\sigma^2 d}$ | $\times$ | ✓ |
| Averaged SGD [1, 9] | $\mathcal{O}\left(\frac{R^4}{\lambda_{\min}^2 n^2} \cdot \text{init} + \frac{\sigma^2 d}{n}\right)$ | $nd$ | $n$ | ✓ | $\times$ |
| Streaming SVRG (with initial error oracle) [12] | $\mathcal{O}\left(\exp\left(-\frac{n\lambda_{\min}(\mathbf{H})}{R^2}\right) \cdot \text{init}\right) + \frac{\sigma^2 d}{n}$ | $nd$ | $(\frac{R^2}{\lambda_{\min}(\mathbf{H})}) \cdot \log \frac{n \cdot \text{init}}{\sigma^2 d}$ | ✓ | ✓ |
| Algorithm 2 | $\mathcal{O}\left(\left(\frac{R^2 t}{\|\mathbf{H}\|_2 n}\right)^{\frac{t}{\kappa \log(d\kappa)}} \cdot \text{init} + \frac{\sigma^2 d}{n}\right)$ | $nd$ | $t \log\left(\frac{n\|\mathbf{H}\|_2}{tR^2}\right)$ | ✓ | ✓ |
| Algorithm 2 with initial error oracle | $\mathcal{O}\left(\exp\left(-\frac{n\lambda_{\min}(\mathbf{H})}{R^2 \cdot \log(d\kappa)}\right) \cdot \text{init} + \frac{\sigma^2 d}{n}\right)$ | $nd$ | $\kappa \log(d\kappa) \log \frac{n \cdot \text{init}}{\sigma^2 d}$ | ✓ | ✓ |

Table 1: Comparison of Algorithm 2 with existing algorithms given $n$ samples for LSR, with init $= R(\mathbf{w}_0) - R(\mathbf{w}^*)$. Note that the bounds have been provided assuming realizable noise; for the algorithms with support for agnostic noise, these can be appropriately modified. Refer to Section 3 for the definitions of all quantities. We do not consider accelerated variants in this table. Note that the accelerated variants serve to improve running times of the offline algorithms. Furthermore, offline algorithms possess the advantage that they do not require the knowledge of $R^2$. The quantity $t$ in the results for Algorithm 2 can be chosen arbitrarily and trades off the rate of decay in initial error with the depth of the algorithm. Finally, we note that streaming SVRG [12] does not conform to the standard first order oracle model [14].

providing a new sharp analysis in order to bound the resulting operators. In particular, straightforward operator norm bounds of these matrix operators suffices to show convergence of SGD updates, but turn out to be pretty loose bounds. In order to obtain stronger bounds, this paper employs a much more fine grained analysis that directly bounds the trace of these operators applied to the relevant matrices, which results in obtaining much tighter bounds on the performance of mini-batched tail-averaged SGD while simultaneously shedding light on the role of agnostic noise on the behavior of SGD.

While this paper's results focus on the case of strongly convex streaming least square regression, we believe that our techniques extend much more broadly. We hope that this paper will serve as the basis for future work on analyzing SGD and improving the parallelization of large scale algorithms for machine learning.

**Paper organization**: A brief overview of related works is presented in Section 2. Then, the problem setup and notations are introduced in Section 3. The main results of this paper as well as the mini-batched tail averaging SGD algorithm are presented in Section 5. In Section 6, simulation results are presented using a synthetic dataset, and these results corroborate well with the main results on the mini-batch thresholds established by the paper as well as the advantages of mini-batch tail averaged SGD algorithm (i.e. algorithm 1). The proofs of all the claims and theorems are provided in the appendix.

## 2 Related work

We split the work related to this paper's contributions into several algorithm classes:

**Stochastic Approximation:** Seminal works of Ruppert [15], and Polyak and Juditsky [8] showed that averaging iterates of SGD applied to stochastic optimization problems like stochastic least squares regression achieves asymptotically optimal statistical rates. A more recent line of work [16, 9, 10, 11, 17, 18, 1] has provided rigorous non-asymptotic analysis (with explicit constants) of SGD for stochastic approximation problems including linear and logistic regression. Equipped with averaging, these algorithms achieve (up to constants) minimax optimal statistical rates. Note that all these results assume that the noise is *realizable* when stating the generalization error bounds. We are unaware of any results that provide sharp non-asymptotic analysis of SGD and the related step size issues in the general agnostic noise setting. Several specialized

stochastic algorithms such as SVRG [19], SDCA [13] and SAGA [20] have been proposed for the agnostic setting, among which, only the streaming version of SVRG [12] is known to provide minimax optimal statistical rates and is highly parallelizable. Note that offline SVRG [19] is parallelizable as well, but is not streaming and incurs additional $\log$ factors without carefully employing doubling tricks. Note that streaming SVRG [12] does not take advantage of minibatching and its depth depends on a stronger fourth moment condition on the input whereas our mini-batching bounds depend only on the condition number of the covariance matrix. Moreover, our constant-step size SGD algorithm is arguably simpler than streaming SVRG and can be run in a more general setting requiring less information from a single sample. In particular, SGD based algorithms conform to the standard first order oracle model [14], whereas the streaming SVRG [12] does not, since it requires gradient information at two different points from a single sample.

Another key aspect of this paper's main result is a *precise* characterization of the effects of *mini-batching*, which is also studied extensively in this literature [21, 22, 23, 6, 24]. Previous results on mini-batching within the context of stochastic approximation yielded results that suggest increasing minibatch sizes tend to hurt the convergence rate [22, 21]. [6, 24, 21, 23] study mini-batching for various algorithms and show faster decay in the bias term but do not show optimal statistical error rates. Moreover, *all* of these results claim speed-up from mini-batching by comparing the *upper bounds* on the training/generalization error. In contrast, this paper's result for mini-batching compares the exact generalization error for different batch sizes and demonstrates an improvement with larger sizes.

The other scheme for parallelizing stochastic approximation is through *model averaging* [2, 3, 25, 26]. These schemes are communication efficient since they average the final output of independent runs of the algorithm, while being minimax optimal. However, as noted in [26], these schemes do not yield improvements to the bias and it is unclear how these results relate to this work's depth bounds. Other efforts [25, 2, 3] provide guarantees on upperbounds of the error, with only [25] guaranteeing minimax optimality.

**Non-Stochastic Approximation:** In order to utilize the recent developments in computing frameworks, there has been a number of efforts in developing distributed/parallel machine learning algorithms that utilize multi-core [4, 5, 21, 22, 27, 24] and multi-machine [2, 3, 6, 23, 7] computing frameworks. As in the stochastic case, parallelization speedups are argued only in terms of upper bounds of training/generalization error. On the other hand, a distinctive feature of this paper's mini-batching result is that the prescribed parallelization speedups are on the *exact* generalization error on a *per problem basis*, and do not rely on providing linear speedups on weaker upper bounds. Moreover, none of the existing parallelization results characterize regimes where there are unconditional linear speedups on the overall generalization error while maintaining minimax optimal statistical error rates. Indeed, model averaging results are known to be minimax optimal [25, 26], however, the bias term is unimprovable using these schemes, as noted by [26].

**Asynchronous Stochastic Approximation:** There exists a recent line of work [5, 28, 29] that provides linear speedups of SGD in the context of asynchronous stochastic optimization. However, in contrast to our work, these efforts deal with the much harder case of asynchronous optimization, and moreover, these results provide speedups on upper bounds on the training/generalization error while relying on specific input structures such as hard sparsity. Moreover, while [29] proves statistical optimality in an asymptotic sense, none of the other efforts claim statistical optimality.

# 3 Preliminaries

**Problem formulation**: Let $\mathcal{D}$ be a distribution on tuples $(\mathbf{x}, y) \in \mathbb{R}^d \times \mathbb{R}$. Let $R : \mathbb{R}^d \to \mathbb{R}$ be defined as the expected square loss over the pairs $(\mathbf{x}, y)$:

$$R(\mathbf{w}) = \frac{1}{2}\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}\left(y - \langle \mathbf{w}, \mathbf{x}\rangle\right)^2 \ \forall\, \mathbf{w} \in \mathbb{R}^d. \tag{1}$$

Our goal is to compute the predictor $\mathbf{w}^* = \arg\min_{\mathbf{w}} R(\mathbf{w})$. In this paper, we provide the maiden sharpest analysis of SGD with batch sizes $\geq 1$, applied to the general agnostic case of LSR (also referred to as the unrealizable case) where $y \stackrel{\text{def}}{=} \langle \mathbf{w}^*, \mathbf{x}\rangle + \epsilon$ where $\epsilon \in \mathbb{R}$ is the noise.

**Notation**: We use boldface small letters ($\mathbf{x}, \mathbf{w}$ etc.) for vectors and boldface capital letters ($\mathbf{A}, \mathbf{H}$ etc.) for matrices. We use $\otimes$ to denote the outer product of two vectors or matrices. $\mathcal{S}(d)$ denotes the space of real symmetric $d \times d$ matrices and $\mathfrak{M}(d)$ denotes the space of linear transformations from $\mathbb{R}^{d\times d}$ to $\mathbb{R}^{d\times d}$. Every element in $\mathfrak{M}(d)$ is a $d \times d \times d \times d$ tensor, which is represented using normal font scripted letters

**Algorithm 1** Minibatch-TailAveraging-SGD

**Input:** Initial point $\mathbf{w}_0$, stepsize $\gamma$, minibatch size $b$, initial iterations $s$, total number of samples $m$.

**for** $t = 1, 2, .., s$ **do**
  Sample "$b$" tuples $\{(x_{ti}, y_{ti})\}_{i=1}^{b} \sim \mathcal{D}^b$
  $\mathbf{w}_t \leftarrow \mathbf{w}_{t-1} - \frac{\gamma}{b}\sum_{i=1}^{b}(\langle \mathbf{w}_{t-1}, x_{ti}\rangle - y_{ti})x_{ti}$
$\bar{\mathbf{w}} = 0$
**for** $t = s+1, s+2, .., \frac{m}{b}$ **do**
  Sample "$b$" tuples $\{(x_{ti}, y_{ti})\}_{i=1}^{b} \sim \mathcal{D}^b$
  $\mathbf{w}_t \leftarrow \mathbf{w}_{t-1} - \frac{\gamma}{b}\sum_{i=1}^{b}(\langle \mathbf{w}_{t-1}, x_{ti}\rangle - y_{ti})x_{ti}$
  $\bar{\mathbf{w}} = \frac{1}{t-s}\mathbf{w}_t + \frac{t-s-1}{t-s}\bar{\mathbf{w}}$

**Output:** $\bar{\mathbf{w}}$

($\mathcal{M}, \mathcal{T}$ etc.). A tensor $\mathcal{M}$ applied to a matrix $\mathbf{A}$ gives a matrix $\mathcal{M}\mathbf{A}$ such that its $ij^{\text{th}}$ element is given by $(\mathcal{M}\mathbf{A})_{ij} = \sum_{k,\ell\in[d]} \mathcal{M}_{ik\ell j}\mathbf{A}_{k\ell}$. Any $d \times d$ matrix $\mathbf{A}$ can be lifted to $\mathcal{A}_{\mathcal{L}}, \mathcal{A}_{\mathcal{R}} \in \mathfrak{M}(d)$, which represents the left/right multiplication operator of the matrix $\mathbf{A}$ respectively. More concretely, $\mathcal{A}_{\mathcal{L}}, \mathcal{A}_{\mathcal{R}} \in \mathfrak{M}(d)$ are defined so that $\mathcal{A}_{\mathcal{L}}\mathbf{W} = \mathbf{A}\mathbf{W}$ and $\mathcal{A}_{\mathcal{R}}\mathbf{W} = \mathbf{W}\mathbf{A}$ $\forall\ \mathbf{W} \in \mathbb{R}^{d\times d}$. We let $\mathcal{I} \in \mathfrak{M}(d)$ represent the identity map.

We let $\|\cdot\|$ or $\|\cdot\|_2$ denote $\ell_2$ norm for vectors and spectral or operator norm for matrices and tensors. More concretely, we define

$$\|\mathbf{A}\| \stackrel{\text{def}}{=} \max_{\mathbf{x}\in\mathbb{R}^d} \frac{\|\mathbf{A}\mathbf{x}\|}{\|\mathbf{x}\|}\ \forall\ \mathbf{A} \in \mathbb{R}^{d\times d},\ \text{and}\ \|\mathcal{T}\| \stackrel{\text{def}}{=} \max_{\mathbf{A}\in\mathcal{S}(d)} \frac{\|\mathcal{T}\mathbf{A}\|_F}{\|\mathbf{A}\|_F}\ \forall\ \mathcal{T} \in \mathfrak{M}(d),$$

where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix. When we refer to the largest/smallest eigenvalues $\lambda_{\max}(\mathcal{S})$ or $\lambda_{\min}(\mathcal{S})$ of a tensor $\mathcal{S}$, this is obtained by viewing the operator $\mathcal{S}$ as a member of $\mathfrak{M}(d)$.

**Relevant quantities**: We denote by $\mathbf{H}$ and $\mathcal{M}$ the population second moment matrix and fourth moment tensor of the covariates $\mathbf{x}$, i.e.:

$$\mathbf{H} \stackrel{\text{def}}{=} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mathbf{x}\otimes\mathbf{x}],\ \ \mathcal{M} \stackrel{\text{def}}{=} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mathbf{x}\otimes\mathbf{x}\otimes\mathbf{x}\otimes\mathbf{x}]$$

We assume that both $\mathbf{H}$ and $\mathcal{M}$ are finite and are related by $\mathcal{M}\mathbf{I} \preceq R^2\mathbf{H}$. This in particular implies that $\mathbb{E}\left[\|\mathbf{x}\|^2\right] \le R^2$. We further assume that $\mathbf{H}$ is non-singular, with $\lambda_{\min}(\mathbf{H}) > 0$ denoting the smallest eigenvalue of $\mathbf{H}$. This makes $R(\mathbf{w})$ strongly convex and implies the existence of a unique minimizer $\mathbf{w}^*$ of the objective $R(\mathbf{w})$. The condition number $\kappa$ of the population covariance matrix $\mathbf{H}$ is defined to be the ratio of the largest to smallest eigenvalues of $\mathbf{H}$, i.e., $\kappa \stackrel{\text{def}}{=} \frac{\lambda_{\max}(\mathbf{H})}{\lambda_{\min}(\mathbf{H})}$. Next, we define the noise covariance matrix $\mathbf{\Sigma}$ as:

$$\mathbf{\Sigma} \stackrel{\text{def}}{=} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}\left[(y - \langle\mathbf{w}^*, \mathbf{x}\rangle)^2 \cdot \mathbf{x}\otimes\mathbf{x}\right]$$

In the realizable case, the noise $\epsilon$ is independent of the covariate $\mathbf{x}$, thus implying $\mathbf{\Sigma} = \sigma^2\mathbf{H}$. Note that $\mathcal{M}$ can also be thought of as an operator in $\mathfrak{M}(d)$ which acts on matrix $\mathbf{A}$ to produce $\mathcal{M}\mathbf{A} = \mathbb{E}\left[(\mathbf{x}^\top\mathbf{A}\mathbf{x})\mathbf{x}\otimes\mathbf{x}\right]$. For a given $\gamma > 0$ and $b \in \mathbb{Z}^+$, a key quantity in our results turns out to be the operator $\mathcal{T}_b^\gamma$ defined as $\mathcal{T}_b^\gamma \stackrel{\text{def}}{=} \mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma}{b}\mathcal{M} - \frac{\gamma(b-1)}{b}\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}$. In most cases we drop $\gamma$ for notational convenience and use $\mathcal{T}_b$ for $\mathcal{T}_b^\gamma$. We let $\mathcal{T} \stackrel{\text{def}}{=} \mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \gamma\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}$ denote the limit of $\mathcal{T}_b$ as $b \to \infty$. For a given $\gamma > 0$, we also use $\mathcal{G}$ to denote the operator $\mathcal{G} \stackrel{\text{def}}{=} {\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma\mathcal{I}$.

**Assumptions:** For the realizable case, the generalization error estimate in Theorem 1 is exact on a per-problem basis if $\mathbb{E}\left[y_i|\mathbf{x}_i\right] = \langle\mathbf{w}^*, \mathbf{x}_i\rangle$ (which is true if we have a constant component in $\mathbf{x}$, as is typically the case in practice). For the agnostic case, we note that this result is exact if $\mathbb{E}\left[x^{(i)}x^{(j)}x^{(k)}\epsilon\right] = 0\ \forall\ (i,j,k) \in \otimes_3\{1, 2, .., d\}$ with $x^{(i)}$ referring to the $i^{th}$ dimension of $\mathbf{x}$. If this is not the case, the generalization error in Theorem 1 are off the exact value by a factor of at most 2 (as noted by [9, 1], owing to the application of Minkowski's inequality).

# 4 Viewing Stochastic Gradient Descent as a Stochastic Process

In this section, we consider a stochastic process view of Stochastic Gradient Descent updates and shed some light on why this view allows us to provide a sharp characterization of SGD as a tool to solve the general LSR

**Algorithm 2** MinibatchDoublingPartialAveragingSGD

**Input:** Initial point $\mathbf{w}_0$, stepsize $\gamma$, initial minibatch size $b$, number of iterations in each epoch $s$, number of samples $n$.

/*Run logarithmic number of epochs where each epoch runs $t$ iterations of minibatch SGD (with out averaging). Double minibatch size after each epoch.*/

**for** $\ell = 1, 2, \cdots, \log \frac{n}{bt} - 1$ **do**

  $b_\ell \leftarrow 2^{\ell-1} b$

  $\mathbf{w}_\ell \leftarrow$ Minibatch-TailAveraging-SGD$(\mathbf{w}_{\ell-1}, \gamma, b_\ell, t-1, t \cdot b_\ell)$

/*For the last epoch, run tail averaged minibatch SGD with initial point $\mathbf{w}_t$, stepsize $\gamma$, minibatch size $2^{\log \frac{n}{bt}-1} \cdot b = n/2t$, number of initial iterations $t/2$ and number of samples $n/2$.*/

$\overline{\mathbf{w}} \leftarrow$ Minibatch-TailAveraging-SGD$(\mathbf{w}_s, \gamma, n/2t, t/2, n/2)$

**Output:** $\overline{\mathbf{w}}$

problem. In particular, let us consider the mini-batch SGD update that moves iterate $\mathbf{w}_{t-1}$ to $\mathbf{w}_t$:

$$\mathbf{w}_t = \mathbf{w}_{t-1} - \frac{\gamma}{b} \sum_{i=1}^{b} (\langle \mathbf{w}_{t-1}, \mathbf{x}_{ti} \rangle - y_{ti}) \mathbf{x}_{ti}$$

By defining $\eta_t \stackrel{\text{def}}{=} \mathbf{w}_t - \mathbf{w}^*$, we write the recursion relating $\eta_t$ to $\eta_{t-1}$ as:

$$\eta_t = \big(\mathbf{I} - \frac{\gamma}{b} \sum_i \mathbf{x}_{ti} \otimes \mathbf{x}_{ti}\big) \eta_{t-1} + \frac{\gamma}{b} \sum_i \epsilon_{ti} \mathbf{x}_{ti}$$

Now, since we deal with the square loss case, the generalization error is written as:

$$\mathbb{E}\left[R(\mathbf{w}_t)\right] - R(\mathbf{w}^*) = \frac{1}{2} \operatorname{Tr} \mathbf{H}\, \mathbb{E}\left[\eta_t \otimes \eta_t\right]$$

We thus consider $\mathbb{E}\left[\eta_t \otimes \eta_t\right]$, which can be written in terms of $\mathbb{E}\left[\eta_{t-1} \otimes \eta_{t-1}\right]$ as:

$$\mathbb{E}\left[\eta_t \otimes \eta_t\right] = (\mathcal{I} - \gamma \mathcal{T}_b)\, \mathbb{E}\left[\eta_{t-1} \otimes \eta_{t-1}\right] + \gamma^2 \mathbf{\Sigma}$$

**Divergent Learning Rates:** This operator view of mini-batch SGD immediately sheds light on the maximal stepsizes $\gamma^{div}_{b,\max}$ that prevents divergence of the resulting mini-batch SGD updates: in particular, a necessary condition that the set of non-divergent learning rates $\gamma \geq 0$ should satisfy is $\mathcal{T}_b^\gamma \succeq 0$. This condition yields the following generalized eigenvector problem for computing $\gamma^{div}_{b,\max}$:

$$\frac{2}{\gamma^{div}_{b,\max}} = \sup_{\mathbf{W} \in \mathcal{S}(d)} \frac{\langle \mathbf{W}, \mathcal{M}\mathbf{W} \rangle + (b-1) \cdot \operatorname{Tr} \mathbf{WHWH}}{b \cdot \operatorname{Tr} \mathbf{WHW}}$$

This characterization generalizes the divergent stepsize characterization of [1] for batch sizes $> 1$. We note that such a characterization automatically sheds light on how the maximal divergent learning rates interpolate batch sizes of 1 (which is $\leq \frac{2}{\operatorname{Tr} \mathbf{H}}$ [1]) to the batch gradient descent learning rate (setting $b$ to $\infty$), which turns out to be $\frac{2}{\lambda_{\max}(\mathbf{H})}$. A property of $\gamma^{div}_{b,\max}$ that is worth noting is that it does not depend on properties of the noise covariance $\mathbf{\Sigma}$, and depends solely on the statistics of the covariate $\mathbf{x}$.

We note that in this paper, our interest does not lie in the non-divergent stepsizes $0 \leq \gamma \leq \gamma^{div}_{b,\max}$, but in the set of (maximal) step sizes $0 \leq \gamma \leq \gamma_{b,\max}$ ($< \gamma^{div}_{b,\max}$) that are sufficient to guarantee statistically optimal error rates of $\mathcal{O}(\frac{1}{n} \operatorname{Tr} \mathbf{H}^{-1} \mathbf{\Sigma})$. We show that our characterization of $\gamma_{b,\max}$ is indeed tight in the sense that there are problem instances, where $\gamma_{b,\max}$ and $\gamma^{div}_{b,\max}$ are off by a factor of $d$ and $\gamma_{b,\max}$ is still the largest step size which provides statistically optimal error rates. A critical property of $\gamma_{b,\max}$ is that it is dependent on the properties of the noise covariance $\mathbf{\Sigma}$. We return to the characterization of $\gamma_{b,\max}$ in section 5.4.

# 5 Main Results

This section presents the main Algorithms analyzed by this paper and their associated guarantees. In particular, the standard mini-batch tail-averaged SGD Algorithm for the streaming LSR problem is presented

as Algorithm 1. Theorem 1 presents an exact characterization of the error achieved by Algorithm 1. Theorem 2 establishes a scalar estimate of the generalization error while dealing with the general agnostic case of LSR, in order to better illustrate the performance of Algorithm 1. Theorem 2 shows that increasing minibatch sizes (up to some threshold) and tail averaging achieves the same work (number of first order oracle calls to get gradient at a point) and statistical error performance as the standard averaged SGD but decreases the depth significantly, thereby increasing its parallelizability.

Theorem 3 presents the guarantees of Algorithm 2, which is a doubling based Algorithm that decays the initial error at a superpolynomial rate (and a geometric rate in the presence of an initial error oracle) while performing the same work and achieving same statistical accuracy (up to constant factors) compared to standard sequential SGD (with batch size 1). Algorithm 2's depth nearly matches the depth of standard offline gradient descent, while providing the same generalization error (upto constants) as offline gradient descent on *every* problem instance, while being a single pass streaming Algorithm, thus performing substantially lower work compared to the standard offline gradient descent. Refer to table 1 for the associated bounds. The paper's final result is lemma 4, which presents an instance that shows learning rates necessary to retain statistical optimality in the agnostic case could be much lower than ones required by realizable case.

## 5.1 Characterization of SGD with minibatching and tail averaging

This section presents Theorem 1, which provides an exact characterization of the error incurred by mini-batched tail-averaged SGD, as described in Algorithm 1, working with a mini-batch size $b$ and step size $\gamma$. Algorithm 1 consists of two phases. The first phase (lines 1-3) is the burn-in phase that consists of "$s$" iterations where mini-batch SGD updates are performed *without* averaging the iterates. The second phase (lines 5-8) is the tail-averaging phase, where, iterates of mini-batch SGD are averaged. Intuitively, the burn-in phase is used to decrease the bias which is dependent on initial error $\|\mathbf{w}_0 - \mathbf{w}^*\|$ and the averaging phase is to decrease the variance. Note that the bias exhibits a geometric decay during the unaveraged phase.

**Theorem 1.** *For all $b \geq 1$, let $\gamma_{b,max}$ be defined as $\gamma_{b,max} \stackrel{\text{def}}{=} \min\{\frac{2b}{(b-1)\|\mathbf{H}\|_2}, \frac{2b\,\mathrm{Tr}\,(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}{\|(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}\|_2 R^2 d}\}$. Suppose we run Algorithm 1 with initial point $\mathbf{w}_0$, stepsize $0 < \gamma < \gamma_{b,max}$, minibatch size $b$, initial iterations $s$, total number of samples $n$, and $\exists$ universal constants $C_1, C_2$ such that the final error $\overline{\mathbf{w}} - \mathbf{w}^*$ satisfies:*

$$\mathbb{E}\left[(\overline{\mathbf{w}} - \mathbf{w}^*)(\overline{\mathbf{w}} - \mathbf{w}^*)^\top\right] = \underbrace{\frac{1}{(\frac{n}{b}-s)^2\gamma^2}\mathcal{G}\mathcal{T}_b^{-1}(\mathcal{I}-\gamma\mathcal{T}_b)^s\mathbf{E_0}}_{\mathfrak{T}_1} + \underbrace{C_1\cdot\frac{\rho_b^{n/b}}{\gamma\lambda_{min}(\mathbf{H})\left(\frac{n}{b}-s\right)}\|\mathbf{E_0}\|_F\cdot\mathbf{M_1}}_{\mathfrak{T}_2}$$
$$+\underbrace{\frac{1}{n-sb}\left[\mathcal{I}-\frac{1}{\gamma\left(\frac{n}{b}-s\right)}(\mathcal{H}_{\mathcal{L}}{}^{-1}+\mathcal{H}_{\mathcal{R}}{}^{-1})\right]\mathcal{G}\mathcal{T}_b{}^{-1}\mathbf{\Sigma}}_{\mathfrak{T}_3}$$
$$+\underbrace{\frac{2}{\gamma b\left(\frac{n}{b}-s\right)^2}\mathcal{H}_{\mathcal{L}}{}^{-1}\mathcal{H}_{\mathcal{R}}{}^{-1}\mathcal{T}_b{}^{-1}\mathbf{\Sigma}}_{\mathfrak{T}_4} + \underbrace{C_2\cdot\frac{\rho_b^{\frac{n}{b}-s+1}}{(n-bs)\lambda_{min}(\mathbf{H})\lambda_{min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F\cdot\mathbf{M_2}}_{\mathfrak{T}_5}, \tag{2}$$

*where $\mathbf{E_0} \stackrel{\text{def}}{=} (\mathbf{w}_0 - \mathbf{w}^*)(\mathbf{w}_0 - \mathbf{w}^*)^\top$ and $\rho_b \stackrel{\text{def}}{=} \max(\|\mathbf{I} - \gamma\mathbf{H}\|, \|\mathcal{I} - \gamma\mathcal{T}_b\|) < 1$ and $\mathbf{M_1}, \mathbf{M_2}$ are matrices whose frobenius norm is less than or equal to 1.*

Under the fourth moment assumption on the covariates $\mathcal{M}\mathbf{I} \preceq R^2\mathbf{H}$, the stepsize $\gamma_{b,\text{max}}$ is the *maximal* step size we can use with SGD for the general agnostic LSR problem while being able to *guarantee minimax optimal statistical error rates*. We return to this step size characterization after presenting this paper's main results. Lemma 7 (in the appendix) provides details regarding why this learning rate is a maximal; briefly put, this stepsize limit characterizes the regime where the operators that govern the SGD update rules possess favorable contractive properties while guaranteeing minimax optimal statistical rates.

Though (2) looks quite complicated, each of the terms in the equation has a very intuitive conceptual meaning. The terms dependent on $\mathbf{E_0}$ (i.e., $\mathfrak{T}_1$ and $\mathfrak{T}_2$) are the *bias terms* while those depending on $\mathbf{\Sigma}$ (i.e., $\mathfrak{T}_3, \mathfrak{T}_4$ and $\mathfrak{T}_5$) are the *variance terms*. We observe the geometric decay of the bias terms $\mathfrak{T}_1, \mathfrak{T}_2$ during the "$s$" initial iterations when the iterates are not averaged. The term $\mathfrak{T}_3$ (which is the leading order variance term) contains a coefficient $\frac{1}{n-sb}$, which is the number of samples that have been used in the averaging phase.

## 5.2 Benefits of minibatching and tail-averaging

This section presents a means to combine minibatching and tail averaging in order to obtain better performance than standard averaged SGD for the general agnostic case of LSR. We note that these results naturally interpolate to results for the realizable case by setting $\mathbf{\Sigma} = \sigma^2\mathbf{H}$. Tail averaging allows a burn-in phase when iterates are not averaged, thereby decaying the initial error geometrically during this phase, while minibatching decreases the depth (i.e. enables parallelizing SGD) while performing the same amount of work (up to constants). By the same amount of work, we mean that after using the same number of samples $n$, the generalization error is the *same* as SGD with a batch-size 1 (up to constant factors): i.e., we obtain near linear parallelization speedups via mini-batching.

**Theorem 2.** *Consider the general agnostic case of LSR in the setting of Theorem 1. If a minibatch size of* $b \geq b_{thresh} \stackrel{\text{def}}{=} \frac{R^2 d \left\| (\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma} \right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}$ *and stepsize* $\gamma = \frac{1}{3\|\mathbf{H}\|_2} \leq \frac{\gamma_{b,max}}{6}$ *is chosen, we obtain the following generalization error bound on the performance of Algorithm 1:*

$$
\begin{aligned}
R(\overline{\mathbf{w}}) - R(\mathbf{w}^*) &= \operatorname{Tr}\left(\mathbf{H}\mathbb{E}\left[(\overline{\mathbf{w}} - \mathbf{w}^*)(\overline{\mathbf{w}} - \mathbf{w}^*)^\top\right]\right) \\
&\leq \underbrace{\frac{36\sqrt{d}\kappa^2}{\left(\frac{n}{b} - s\right)^2} \exp\left(-\frac{s}{3\kappa}\right) \{R(\mathbf{w}_0) - R(\mathbf{w}^*)\}}_{\mathfrak{T}_1} + \underbrace{\exp\left(-\frac{n}{3b\kappa}\right) \frac{3\kappa^2\sqrt{d}}{\frac{n}{b} - s} \{R(\mathbf{w}_0) - R(\mathbf{w}^*)\}}_{\mathfrak{T}_2} \\
&+ \underbrace{\frac{3 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{b(\frac{n}{b} - s)}}_{\mathfrak{T}_3} + \underbrace{\frac{9\kappa \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{b\left(\frac{n}{b} - s\right)^2}}_{\mathfrak{T}_4} + \underbrace{\frac{2\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{b(\frac{n}{b} - s)} \exp\left(-\frac{(\frac{n}{b} - s + 1)}{3\kappa}\right)}_{\mathfrak{T}_5}.
\end{aligned} \tag{3}
$$

**Tail averaging** enables the bias terms $\mathfrak{T}_1$ and $\mathfrak{T}_2$ to decay exponentially with number of initial iterations $s$, which is burn-in phase when iterates of SGD are not averaged. Compared to standard SGD which reduces the initial error $R(\mathbf{w}_0)$ as $\mathcal{O}(\frac{1}{s^2})$ [1], Algorithm 1 reduces the dependence on $R(\mathbf{w}_0) - R(\mathbf{w}^*)$ geometrically as $\exp\left(-\frac{s}{\kappa}\right)$. We note that this contraction factor is reminiscent of the contraction factor that is observed with deterministic batch gradient descent. This paper's analysis of minibatching thus characterizes how the behavior of mini-batch SGD smoothly interpolates from vanilla SGD (with a batch size 1) to batch gradient descent with mini-batch sizes of $\mathcal{O}\left(\frac{R^2 d \left\| (\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma} \right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}((\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma})}\right)$.

**Minibatching** with a batchsize $b = b_{\text{thresh}}$ allows us reduce the depth (serial work) while achieving (up to constants) the statistically optimal rate of $\frac{\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}$. Note that while standard averaged SGD has a depth of $n$, Algorithm 1 has a depth of $n/b$. We note that this reduction in depth does not incur an extra overhead in terms of work (the extra work overhead implies that the Algorithm is wasteful of samples, or in other words, we no longer achieve linear parallelization speedups as a function of increasing the mini-batch size). We refer the reader to **Table 2**, which compares the work, depth and accuracy achieved by minibatching and tail-averaging as compared to averaged SGD with $n$ samples where minibatch size $b = b_{\text{thresh}} = \frac{R^2 d \left\| (\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma} \right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}$ and number of initial steps $s = \frac{n}{2b}$.

## 5.3 (Nearly) Matching the depth of Batch Gradient Descent

The result of section 5.2 establishes a scalar generalization error bound of Algorithm 1 for the general agnostic case of LSR and showed that the depth (number of sequential updates in our algorithm) is decreased to $n/b$. This section builds upon this result to present a simple and intuitive doubling based streaming algorithm that works in epochs and processes a total of $n/2$ points. In each epoch, the minibatch size is increased by a factor of 2 while thus applying Algorithm 1 with twice as many samples as the previous epoch. After running over $n/2$ samples using this epoch based Algorithm, we run Algorithm 1 with the remaining $n/2$ points. Intuitively each of these epochs decreases the bias linearly and then resets the bias of the next epoch as the total risk of previous epoch, and then again decays the bias. The final application of the tail-averaging algorithm is to ensure that the variance is small.

| | Final error | Runtime | Depth |
|---|---|---|---|
| Sequential averaged SGD ($b = 1$) | $\mathcal{O}\left(\frac{\sqrt{d}R^4}{\lambda_{\min}^2 n^2}\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{\sigma^2 d}{n}\right)$ | $\mathcal{O}(nd)$ | $n$ |
| Minibatch-TailAveraging-SGD (Algorithm 1) | $\mathcal{O}\left(\frac{\sqrt{d}R^4}{\lambda_{\min}^2 n^2}\exp\left(-\frac{n\lambda_{\min}}{R^2}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{\sigma^2 d}{n}\right)$ | $\mathcal{O}(nd)$ | $\frac{n}{b_{\text{thresh}}}$ |

Table 2: Comparison of minibatching and tail averaging in Algorithm 1 with with averaged SGD, both using $n$ samples (assuming realizable noise; the agnostic noise bound for Algorithm 1 can be derived similarly). Algorithm 1 uses minibatch size $b_{\text{thresh}} = \frac{R^2 d\left\|(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}\right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}} = \frac{R^2}{\|\mathbf{H}\|_2}$, $\gamma = \frac{1}{3\|\mathbf{H}\|_2}$ and $s = \frac{n}{2b}$ while averaged SGD uses the $n$ samples sequentially. We assume that $n > \frac{R^2}{\lambda_{\min}}$ which is the minimum number of samples to make any progress [12].

The next theorem formalizes this intuition and shows that Algorithm 2 indeed improves the depth exponentially from $n/b_{\text{thresh}}$ to $\mathcal{O}\left(\kappa\log(d\kappa)\log(\frac{n\{R(\mathbf{w}_0)-R(\mathbf{w}^*)\}}{\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}})\right)$ in the presence of an error oracle that provides us with the initial excess risk $R(\mathbf{w}_0) - R(\mathbf{w}^*)$ and the noise level $\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}$.

**Theorem 3.** *Consider the general Agnostic case of LSR. Suppose in MinibatchDoublingPartialAveragingSGD (Algorithm 2), we use initial minibatchsize of* $b = \frac{R^2 d\left\|(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}\right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}$, *stepsize* $\gamma = \frac{1}{3\|\mathbf{H}\|_2} \leq \frac{\gamma_{b,max}}{6}$ *and number of iterations in each epoch being* $t \geq 16\kappa\log(d\kappa)$. *We obtain the following bound on the excess risk of* $\overline{\mathbf{w}}$*:*

$$\mathbb{E}\left[R(\overline{\mathbf{w}})\right] - R(\mathbf{w}^*) \leq \left(\frac{2bt}{n}\right)^{\frac{t}{3\kappa\log(d\kappa)}}\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + 40\frac{\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}.$$

**Remarks**: The final error again has two parts – the bias term that depends on the initial error $R(\mathbf{w}_0)$ and the variance term that depends on the statistical noise $\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}$.

Algorithm 2 decays the bias at a superpolynomial rate by choosing $t$ large enough. If Algorithm 2 has access to an initial error oracle that provides $R(\mathbf{w}_0) - R(\mathbf{w}^*)$, we can run Algorithm 2 with a batch size $b_{\text{thresh}}$ until the excess risk drops to the noise level $\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}$ and subsequently begin doubling the batch size. Such an algorithm indeed gives geometric convergence with a generalization error bound as:

$$\mathbb{E}\left[R(\overline{\mathbf{w}})\right] - R(\mathbf{w}^*) \leq \exp\left(-(\frac{n\lambda_{\min}}{R^2\cdot\log(d\kappa)})\cdot\frac{\operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}{d\left\|(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}\right\|_2}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + 40\frac{\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n},$$

with a depth of $\mathcal{O}\left(\kappa\log(d\kappa)\log\frac{n\{R(\mathbf{w}_0)-R(\mathbf{w}^*)\}}{\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}\right)$. The proof of this claim follows relatively straightforwardly from the proof of Theorem 3. We note that this depth nearly matches (up to $\log$ factors), the depth of standard offline gradient descent. Our final observation is that the statistical error decays at a rate of $\mathcal{O}\left(\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}/n\right)$ which is statistically optimal up to constant factors.

### 5.4 Tradeoffs between mini-batch sizes, learning rates and the role of agnostic noise

While existing sharp non-asymptotic analyses of SGD with large learning rates [9, 10, 17, 1] focus on the realizable case (i.e. the case with independent noise, so that $\mathbf{\Sigma} = \sigma^2\mathbf{H}$), this paper provides results in the agnostic noise case of LSR that point towards fundamental differences in the behavior of SGD between the realizable and agnostic noise setting. Moreover, this paper's results in general agnostic noise case naturally specialize to existing results in the realizable case with batch size 1. Before discussing these differences, we first understand mini-batching and related learning rate issues better. We start by writing out the maximal learning rate that permits statistical optimality:

$$\gamma_{b,\max} = \min\left\{\frac{2b}{(b-1)\|\mathbf{H}\|_2}, \frac{2b\operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}{\|(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}\|_2 R^2 d}\right\}$$

We note that the maximal learning rates $\gamma_{b,\max}$ tends to increase *linearly* as a function of the batch size $b$, until the point that the batch sizes hit a problem dependent limit $b_{\text{thresh}} = \frac{R^2 d \left\| (\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma} \right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-1}\mathbf{\Sigma}}$; this is precisely the range of mini-batch sizes that yield a *provably near-linear parallelization speedup* over standard SGD with batch size $1$, while *retaining statistical optimality*. Moreover, when the batch sizes hit $b_{\text{thresh}}$, the first term in the definition of $\gamma_{b,\max}$ becomes active, and we begin using learning rates that mimic batch gradient descent, i.e., with $\gamma_{b,\max} = \mathcal{O}(\frac{1}{\lambda_{\max}(\mathbf{H})})$. Beyond this mini-batching threshold $b_{\text{thresh}}$, mini-batching is inefficient in that while the variance term still improves as the batch size increases, the rate of decay of the bias does not show a linear improvement as a function of the batch size (owing to the lack of serial updates), thus leading to work inefficiency.

Next, we note that in the realizable case, $\gamma_{b,\max} = \min\{\frac{2b}{(b-1)\|\mathbf{H}\|_2}, \frac{2b}{R^2}\}$. This fact, coupled with the case of SGD with batch size $1$ (as considered in [9, 10, 17]) provides $\gamma_{1,\max} = \frac{2}{R^2}$, which matches the best learning rates known to provide statistical optimality. The corresponding mini-batching threshold for the *realizable* case is $b_{\text{thresh}} = \frac{R^2}{\|\mathbf{H}\|_2}$.

This paper's final result deals with shedding light on some fundamental differences in the behavior of SGD when dealing with the realizable case as opposed to the agnostic noise case. In particular, the following lemma presents an instance working with agnostic noise case, wherein, SGD *cannot* employ the large learning rates that typically are used in the realizable case while *retaining statistical optimality*; the proof of this lemma can be found in the appendix. Note that this behavior is in stark contrast to algorithms such as streaming SVRG [12], which work with the same large learning rates in the agnostic case as in the realizable case, while guaranteeing statistically optimal rates.

**Lemma 4.** *Consider a Streaming LSR example with Gaussian covariates with a diagonal second moment matrix defined by:*

$$\mathbf{H}_{ii} = \begin{cases} 1 & \text{if } i = 1 \\ 1/d & \text{if } i > 1 \end{cases}$$

*Further, let the noise covariance matrix $\mathbf{\Sigma}$ be diagonal as well, with the following entries:*

$$\mathbf{\Sigma}_{ii} = \begin{cases} 1 & \text{if } i = 1 \\ 1/(d-1)d & \text{if } i > 1 \end{cases}$$

*For this problem instance, $\gamma_{1,max} \leq \frac{4}{(d+2)(1+\frac{1}{d})}$ is necessary for retaining statistical optimality, while the realizable version of this problem permits a maximal learning rate $\leq \frac{d}{(d+2)(1+\frac{1}{d})}$, thus implying an $\mathcal{O}(d)$ separation in learning rates between the realizable and agnostic case.*

The proof of lemma 4 can be found in the appendix in section A.4. As a final remark, note that in the realizable case, we are unaware of whether there is a separation between $\mathcal{O}(\frac{1}{R^2})$ learning rate and the divergent step size schema as introduced by [1]. Furthermore, it is unclear whether it is possible if we can prove statistical optimality when working with step sizes that are larger than $\mathcal{O}(\frac{1}{R^2})$. However, we believe this question is relatively of less consequence given the separation that exists between the realizable and agnostic case, where a *much smaller* learning rate is *necessary* for obtaining statistically optimal rates.

# 6 Experimental simulations

We conduct experiments using a synthetic example to illustrate the implications of our theoretical results. The data is sampled from a $50-$ dimensional Gaussian with eigenvalues decaying as $\{\frac{1}{k}\}_{k=1}^{50}$ (condition number $\kappa = 50$), and the variance $\sigma^2$ of the noise is $0.01$. In this case, our estimated batch size according to Theorem 2 is $b_{\text{thresh}} = 11$. In all our plots, the x-axis is the log of number of iterations (at each iteration, we use "$b$" samples in the SGD update rule), and y-axis is the log of the generalization error. We assume that each iteration takes constant time for all batch sizes; this is just in order to present evidence regarding the tightness of our mini-batching characterization of the regime of mini-batch sizes where we obtain problem dependent linear speedups while achieving the same generalization error as SGD with mini-batch size of $1$.

We first consider the effect of mini-batching with batch sizes of $1$, $4$, $b_{\text{thresh}} = 11$, $2 * b_{\text{thresh}} = 22$ and $d = 50$. In figure 1, we observe the generalization error curves obtained by averaging $100$ runs of the Algorithm, each

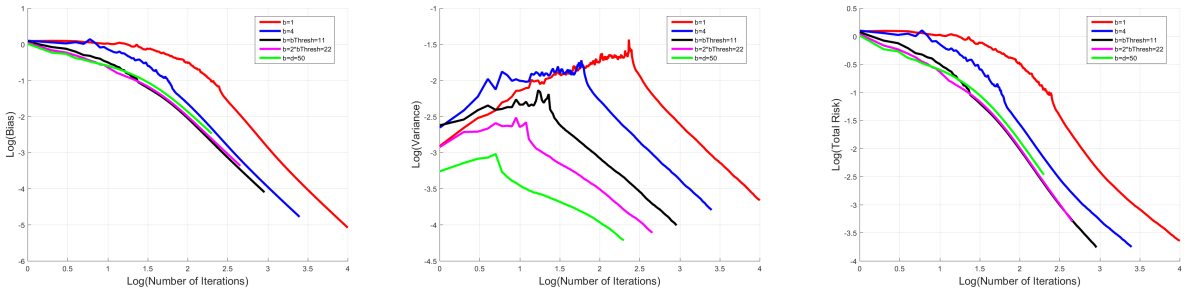


(a) Bias term (b) Variance term (c) Total Risk

Figure 1: Effect of increasing batch sizes on the Algorithm's generalization error. As expected, variance decreases monotonically with increasing batch size. But the bias term shows that the rate of decay increases till the optimal $b_{thresh}$. Total risk shows that with optimal batch size, our method indeed obtains the optimal statistical rate but using smaller number of iterations compared to sub-optimal batch sizes.

run utilizing the same set of samples for all mini-batch sizes. Averaging begins after observing a fixed number of samples (set as $5\kappa$). We see that the rate of bias decay (figure 1a) increases until reaching a mini-batch size of $b_{\text{thresh}}$, saturating thereafter; this implies we are wasteful in terms of sample size. As expected, the rate of decay of variance (figure 1b) is monotonic as a function of mini-batch size. Finally, the overall generalization error (figure 1c) shows the tightness of our mini-batching characterization. This is because, with a batch size of $b_{\text{thresh}}$, we obtain a generalization error that is the *same* as using batch size of $1$ with the number of (serial) iterations that is an order of magnitude smaller. Subsequently, we note that larger mini-batch sizes tend to worsen generalization error thus depicting the tightness of our characterization of $b_{\text{thresh}}$.

In the next experiment, we fix the batch size (as $b_{\text{thresh}}$) as well as the number of iterations that we run mini-batch SGD, and thus hold the total number of samples as $10\kappa$. We then consider the effect of when averaging begins; in particular, we consider averaging iterates right from the start, as prescribed by Dèfossez and Bach [1]. We also consider averaging iterates after a quarter/half/three-fourth of total number of iterations, and consider unaveraged SGD as well. In figure 2, we see that the bias (figure 2a) exhibits a geometric decay in the unaveraged phase while switching to an inferior $\mathcal{O}(\frac{1}{t^2})$ rate with averaging. The variance (figure 2b) tends to increase and stabilize at $\mathcal{O}(\frac{\sigma^2}{b_{\text{thresh}}})$ in the absence of averaging, while switching to a $\mathcal{O}(\frac{1}{N})$ decay rate when averaging begins. The overall generalization error (figure 2c) shows the superiority of the scheme where averaging after a burn-in period allows us decay the bias towards the noise level at a geometric rate, following which tail-averaging allows us to decay the variance term, thus providing credence to our theoretical results that averaging after a burn-in time allows us to obtain better generalization error as a function of sample size.

Finally, we consider the effect of step size choice; in particular, we fix the number of samples to be $10\kappa$ and begin averaging after observing $2\kappa$ samples. We employ a mini-batch size of $b_{\text{thresh}} = 11$ and compute the corresponding $\gamma_{b,\max}$. We then run mini-batch SGD with a step size $\gamma = c \cdot \gamma_{b_{\text{thresh}},\max}$ with $c$ varied as $0.25/0.5/0.75/1/1.1$. The results, as observed in figure 3 presents the effect of increasing stepsizes. In particular, for the bias term (figure 3a), the bias is decays progressively faster during the unaveraged phase as the step size is increased, until we hit $\gamma_{b,\max}$, where the behavior tends to be composed of a combination of divergent and non-divergent steps. The trend is very similar with regard to the behavior of the variance term (figure 3b) and the total risk (figure 3c).

# 7 Concluding remarks

The framework established as a part of Theorem 1 can be extended in a very straightforward manner to provide sharp finite sample generalization error bounds for *parameter mixing/model averaging* schemes [2, 3, 25, 26]. Note that parameter mixing/model averaging is a parallelization scheme that can be used in conjunction with mini-batching and tail-averaging. Furthermore, as noted by [26], these schemes do not offer any improvement on the bias part of the error, while improving the variance term owing to the effect of averaging: this suggests

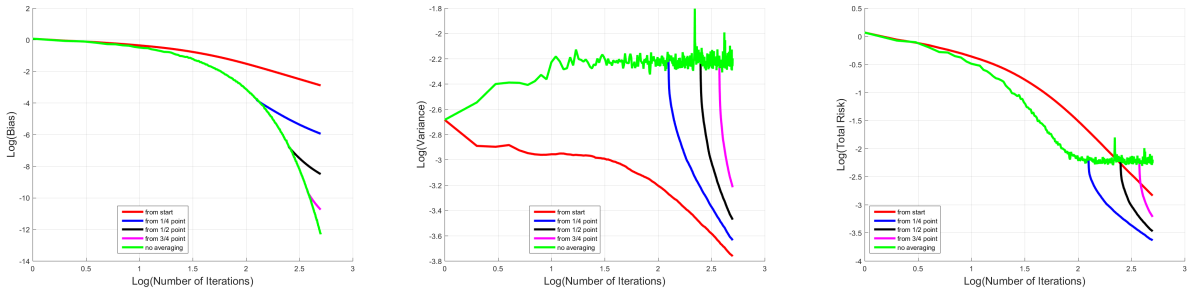


(a) Bias term (b) Variance term (c) Total Risk

Figure 2: Effect of tail-averaging with mini-batch size of $b_{thresh} = 11$.

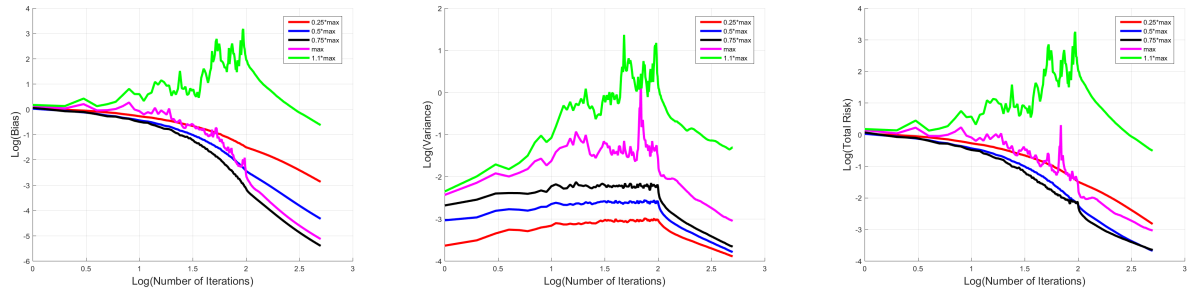


(a) Bias term (b) Variance term (c) Total Risk

Figure 3: Effect of increasing step sizes, i.e. with $\gamma = c \cdot \gamma_{b_{\text{thresh}},\max}$ with $c = 0.25/0.5/0.75/1/1.1$.

the effectiveness of these schemes when the bias error is roughly at the noise level.

While this paper provided the first rigorous analysis of the effect of mini-batching with SGD for the general agnostic case of the streaming LSR problem, we believe the effect of mini-batching holds in much more generality, with corresponding algorithm/problem dependent thresholds on the mini-batch size. In particular, possible future directions could include understanding mini-batching thresholds and effects for stochastic approximation with the Logistic Loss using the framework of [16], more general convex functions [9] and for PCA using the framework of [30]. Moreover, similar mini-batching thresholds could be derived for algorithms other than SGD, such as that of SVRG, by building on the framework established by [12].

# 8 Acknowledgments

Sham Kakade acknowledges funding from the Washington Research Foundation Fund for Innovation in Data-Intensive Discovery. Rahul Kidambi thanks James Saunderson for useful discussions on matrix operator theory.

# References


[1] Alexandre Défossez and Francis R. Bach. Averaged Least-Mean-Squares: Bias-Variance Trade-offs and Optimal Sampling Distributions. In *AISTATS*, volume 38, 2015.

[2] Gideon Mann, Ryan T. McDonald, Mehryar Mohri, Nathan Silberman, and Dan Walker. Efficient Large-Scale Distributed Training of Conditional Maximum Entropy Models. In *NIPS 22*, 2009.

[3] Martin A. Zinkevich, Alex Smola, Markus Weimer, and Lihong Li. Parallelized stochastic gradient descent. In *NIPS 24*, 2011.

[4] Joseph K. Bradley, Aapo Kyrola, Danny Bickson, and Carlos Guestrin. Parallel Coordinate Descent for L1-Regularized Loss Minimization. In *ICML*, 2011.

[5] Feng Niu, Benjamin Recht, Christopher Re, and Stephen J. Wright. Hogwild: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent. In *NIPS 24*, 2011.

[6] Mu Li, Tong Zhang, Yuqiang Chen, and Alexander J. Smola. Efficient mini-batch training for stochastic optimization. In *KDD*, 2014.

[7] Yuchen Zhang and Lin Xiao. DiSCO: Distributed Optimization for Self-Concordant Empirical Loss. In *ICML*, 2015.

[8] B. T. Polyak and A. B. Juditsky. Acceleration of Stochastic Approximation by Averaging. *SIAM J Control Optim*, volume 30, 1992.

[9] Francis Bach and Eric Moulines. Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. In *NIPS 24*, 2011.

[10] Francis R. Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate O(1/n). In *NIPS 26*, 2013.

[11] Francis R. Bach. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. *JMLR*, volume 15, 2014.

[12] Roy Frostig, Rong Ge, Sham M. Kakade, and Aaron Sidford. Competing with the Empirical Risk Minimizer in a Single Pass. In *COLT*, 2015.

[13] Shai Shalev-Shwartz and Tong Zhang. Stochastic Dual Coordinate Ascent Methods for Regularized Loss Minimization. *CoRR*, abs/1209.1873, 2012.

[14] Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 2012.

[15] D. Ruppert. Efficient Estimations from a Slowly Convergent Robbins-Monro Process. *Tech. Report, ORIE, Cornell University*, 1988.

[16] Francis R. Bach. Self-concordant analysis for logistic regression. *Elec. J. of Statistics*, volume 4, 2010.

[17] Aymeric Dieuleveut and Francis Bach. Non-parametric Stochastic Approximation with Large Step sizes. *The Annals of Statistics*, 2015.

[18] Nicolas Flammarion and Francis R. Bach. From Averaging to Acceleration, There is Only a Step-size. In *COLT*, volume 40, 2015.

[19] Rie Johnson and Tong Zhang. Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. In *NIPS 26*, 2013.

[20] Aaron Defazio, Francis R. Bach, and Simon Lacoste-Julien. SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives. In *NIPS 27*, 2014.

[21] Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better Mini-Batch Algorithms via Accelerated Gradient Methods. In *NIPS 24*, 2011.

[22] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal Distributed Online Prediction Using Mini-Batches. *JMLR*, volume 13, 2012.

[23] Shai Shalev-Shwartz and Tong Zhang. Accelerated Mini-Batch Stochastic Dual Coordinate Ascent. In *NIPS 26*, 2013.

[24] Martin Takác, Peter Richtárik, and Nati Srebro. Distributed Mini-Batch SDCA. *CoRR*, abs/1507.08322, 2015.

[25] Yuchen Zhang, John C. Duchi, and Martin Wainwright. Divide and Conquer Ridge Regression: A Distributed Algorithm with Minimax Optimal Rates. *JMLR*, volume 16, 2015.

[26] Jonathan Rosenblatt and Boaz Nadler. On the Optimality of Averaging in Distributed Statistical Learning. *CoRR*, abs/1407.2724, 2014.

[27] Martin Takác, Avleen Singh Bijral, Peter Richtárik, and Nati Srebro. Mini-Batch Primal and Dual Methods for SVMs. In *ICML*, volume 28, 2013.

[28] Horia Mania, Xinghao Pan, Dimitris Papailiopoulos, Benjamin Recht, Kannan Ramchandran, and Michael I. Jordan. Perturbed Iterate Analysis for Asynchronous Stochastic Optimization. *CoRR*, abs/1507.06970, 2015.

[29] John C. Duchi, Sorathan Chaturapruek, and Christopher Ré. Asynchronous stochastic convex optimization. *CoRR*, abs/1508.00882, 2015.

[30] Prateek Jain, Chi Jin, Sham M. Kakade, Praneeth Netrapalli, and Aaron Sidford. Streaming PCA: Matching Matrix Bernstein and Near-Optimal Finite Sample Guarantees for Oja's Algorithm. In *COLT*, 2016.

# A Appendix

**A brief note on the organization of the appendix:**

Theorem 1 establishes the generalization error of Algorithm 1. The proof of Theorem 1 can be found in A.1 and consists of three parts, namely:

- Lemma 7, which rigorously establishes convergence (in section A.1.6) and provides an estimate of the (scalar) contraction factor (in section A.1.7).
- Theorem 5 provides the asymptotic expansion of the finite-sample generalization error of Algorithm 1 on the noiseless problem (also referred to as the bias term), in section A.1.4.
- Theorem 6 provides the asymptotic expansion of the finite-sample generalization error of running Algorithm 1 by initializing it at the solution $\mathbf{w}^*$ (also referred to as the variance term), in section A.1.5.

The proofs of Theorem 2 and Theorem 3 are established in section A.3.

- Theorem 2 decomposes the tensor bound of Theorem 1 and provides a scalar estimate of the generalization performance of the mini-batch tail averaged SGD based Algorithm (Algorithm 1). The proof of Theorem 2 can be found in section A.3.1.
- Theorem 3 provides the generalization error guarantee of the doubling based Algorithm (Algorithm 2). The proof of Theorem 3 can be found in section A.3.2.

The proofs of both Theorem 2 and Theorem 3 utilize the claims established by lemmas in section A.2 along with the tensor bound established in Theorem 1.

Finally, the proof of lemma 4, which describes an instance that yields separation between learning rates required for the realizable and agnostic case is described in section A.4.

**Definition A.1.** We introduce some notations relating the current iterate $\mathbf{w}_t$ and the averaged iterate $\overline{\mathbf{w}}_{t,N}$ to the optimum $\mathbf{w}^*$, namely: $\eta_t \overset{\text{def}}{=} \mathbf{w}_t - \mathbf{w}^*$, and $\bar{\eta}_{t,N} \overset{\text{def}}{=} \overline{\mathbf{w}}_{t,N} - \mathbf{w}^* = \frac{1}{N}\sum_{s=t}^{t+N-1}\eta_s$. Let $\mathbf{L}_t = \mathbb{E}\left[(\mathbf{w}_t - \mathbf{w}^*) \otimes (\mathbf{w}_t - \mathbf{w}^*)\right]$, $\mathbf{L}_{t,N} = \mathbb{E}\left[(\overline{\mathbf{w}}_{t,N} - \mathbf{w}^*) \otimes (\overline{\mathbf{w}}_{t,N} - \mathbf{w}^*)\right]$ denote the covariance of the respective parameter errors. Furthermore, we define the following contraction factors, $\rho_{Tb} \overset{\text{def}}{=} \|\mathbf{I} - \gamma\mathcal{T}_b(\gamma)\|_2$, $\rho_b \overset{\text{def}}{=} \max\{\rho_{Tb}, \|I - \gamma\mathbf{H}\|\}$.

**Theorem 5.** *(Asymptotic Expansion of the error covariance for the bias term) Let $\mathbf{E_0} = \eta_0 \otimes \eta_0$ be the covariance of the initial parameter error. If $0 < \gamma < \gamma_{b,max}$ and if $\epsilon_{t,i} = 0 \; \forall \; i = \{1, ..., b\}, \; t > 0$ (i.e. we deal with the noiseless problem), then, by averaging the iterates of "b" sample mini-batch SGD for $N$ iterations starting from the $s^{th}$ iteration, the covariance term of the error has the following asymptotic expansion:*

$$\mathbf{L}_{s,N}^{bias} = \mathbb{E}[\bar{\eta}_{s,N} \otimes \bar{\eta}_{s,N}] = \frac{1}{N^2\gamma^2}\left[{\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma\mathcal{I}\right]{\mathcal{T}_b}^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s\mathbf{E_0} + C_1 \cdot \frac{\rho_b^{s+N}}{\gamma\lambda_{min}(\mathbf{H})N}\|\mathbf{E_0}\|_F \cdot \mathbf{M}_1 \quad (4)$$

*Where $C_1$ is a universal constant and $\mathbf{M}_1$ is a matrix with Frobenius norm $\leq 1$.*

**Theorem 6.** *(Asymptotic Expansion of the error covariance for the variance term) Let $\mathbf{\Sigma} = \mathbb{E}[\epsilon^2\mathbf{x} \otimes \mathbf{x}]$ and let $\eta_0 = 0$ (i.e. we start at the solution). If $0 < \gamma < \gamma_{b,max}$, then, by averaging the iterates of "b" sample mini-batch SGD for $N$ iterations starting from the $s^{th}$ iteration, the covariance term of the error has the following asymptotic expansion:*

$$\begin{aligned}\mathbf{L}_{t,N}^{variance} = \mathbb{E}[\bar{\eta}_{s,N} \otimes \bar{\eta}_{s,N}] = \frac{1}{bN}\left[\mathcal{I} - \frac{1}{\gamma N}({\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1})\right]\left[(\mathcal{H}_{\mathcal{L}})^{-1} + (\mathcal{H}_{\mathcal{R}})^{-1} - \gamma\mathcal{I}\right]{\mathcal{T}_b}^{-1}\mathbf{\Sigma} \\ + \frac{2}{\gamma bN^2}{\mathcal{H}_{\mathcal{L}}}^{-1}{\mathcal{H}_{\mathcal{R}}}^{-1}{\mathcal{T}_b}^{-1}\mathbf{\Sigma} + C_2 \cdot \frac{\rho_b^{N+1}}{bN\lambda_{min}(\mathbf{H})\lambda_{min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F \cdot \mathbf{M}_2 \quad (5)\end{aligned}$$

*Where $C_2$ is a universal constant and $\mathbf{M}_2$ is a matrix with Frobenius norm $\leq 1$.*

**Lemma 7.** *If $0 < \gamma < \gamma_{b,max}$, then, $\mathcal{T}_b \succ 0$ and $\rho_b < 1$. Furthermore, for dimension $d \geq 2$, the contraction factor $\rho_b$ can be upper bounded as:*

$$\rho_b \leq \max(1 - 2\gamma(1 - \frac{\gamma}{\gamma_{b,\max}})\lambda_{min}(\mathbf{H}), -1 + 2\gamma(1 - \frac{\gamma}{\gamma_{b,max}})\lambda_{max}(\mathbf{H}) + 2(\frac{\gamma}{\gamma_{b,max}})^2, 1 - \gamma\lambda_{min}(\mathbf{H}))$$

*while in dimension $d = 1$, a very similar statement holds, i.e.:*

$$\rho_b \leq \max(1 - 2\gamma(1 - \frac{\gamma}{\gamma_{b,\max}})\lambda_{min}(\mathbf{H}), -1 + 2\gamma(1 - \frac{\gamma}{\gamma_{b,max}})\lambda_{max}(\mathbf{H}) + 2(\frac{\gamma}{\gamma_{b,max}})^2, |1 - \gamma\lambda_{min}(\mathbf{H})|)$$

*Remark* 1. We note that $\gamma \leq \gamma_{b,\max}$ is *sufficient* to ensure Algorithm 1 with constant step sizes achieves statistically optimal rates. Note, our characterization of $\gamma_{b,\max}$ ensures that $\mathcal{T}_b \succeq 0$ (refer to the proof of lemma 7), thus ensuring contraction of all operators that characterize both the burn-in phase as well as averaging phases of the Algorithm.

## A.1 Proofs of Theorems 5, 6 and Lemma 7

We begin by detailing the notation and setup required for proofs of Theorems 5, 6 and lemma 7.

### A.1.1 Deriving and setting up the basic recursion

At each iteration $t$ of Algorithm 1, we are provided with $b$ fresh samples $\{(\mathbf{x}_{ti}, y_{ti})\}_{i=1}^{b}$ drawn i.i.d. from the distribution $\mathcal{D}$. We start by recounting the mini-batch gradient descent update rule that allows us to move from iterate $\mathbf{w}_{t-1}$ to $\mathbf{w}_t$:

$$\mathbf{w}_t = \mathbf{w}_{t-1} - \frac{\gamma}{b}\sum_{i=1}^{b}(\langle \mathbf{w}_{t-1}, \mathbf{x}_{ti}\rangle - y_{ti})\mathbf{x}_{ti}$$

where, $0 < \gamma < \gamma_{b,\max}$ is the constant step size that is set to a value less than the maximum allowed learning rate $\gamma_{b,\max}$. We also recount the definition of $\overline{\mathbf{w}}_{t,N}$ which is the iterate obtained by averaging for $N$ iterations starting from the $t^{th}$ iteration, i.e.,

$$\overline{\mathbf{w}}_{t,N} = \frac{1}{N}\sum_{s=t}^{t+N-1}\mathbf{w}_s$$

Let us first denote the residual error term by $\epsilon_i = y_i - \langle \mathbf{w}^*, \mathbf{x}_i\rangle$. By the first order optimality conditions of $\mathbf{w}^*$, we observe that $\epsilon$ and $\mathbf{x}$ are uncorrelated, i.e, $\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\epsilon\mathbf{x}] = 0$. However, independence of $\epsilon$ and $\mathbf{x}$ is guaranteed only iff $\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[y|\mathbf{x}] = \langle \mathbf{w}^*, \mathbf{x}\rangle$, i.e., in other words, if the noise is well specified. Note that the risk of the iterate $\overline{\mathbf{w}}_{t,N}$ is written as $\bar{R}_{t,N} \overset{\text{def}}{=} R(\bar{w}_{t,N})$. We can now write out the the generalization error as:

$$\mathbb{E}[\bar{R}_{t,N}] - R_* = \frac{1}{2}\operatorname{Tr}\mathbf{H}\,\mathbb{E}\,\bar{\eta}_{t,N}\otimes\bar{\eta}_{t,N} = \frac{1}{2}\mathbb{E}\langle\bar{\eta}_{t,N}, \mathbf{H}\,\bar{\eta}_{t,N}\rangle$$

We now write out the main recursion governing the mini-batch SGD updates in terms of $\eta_.$:

$$\begin{aligned}
\eta_t &= \Big(\mathbf{I} - \frac{\gamma}{b}\sum_{i=1}^{b}\mathbf{x}_{ti}\otimes\mathbf{x}_{ti}\Big)\eta_{t-1} + \frac{\gamma}{b}\sum_{i=1}^{b}\epsilon_{ti}\mathbf{x}_{ti} \\
&= \Big(\mathbf{I} - \frac{\gamma}{b}\sum_{i=1}^{b}\mathbf{x}_{ti}\otimes\mathbf{x}_{ti}\Big)\eta_{t-1} + \frac{\gamma}{b}\sum_{i=1}^{b}\xi_{ti} \\
&= \mathbf{P}_{tb}\eta_{n-1} + \gamma\zeta_{tb} \qquad (6)
\end{aligned}$$

Where, $\mathbf{P}_{tb} \overset{\text{def}}{=} \Big(\mathbf{I} - \frac{\gamma}{b}\sum_{i=1}^{b}\mathbf{x}_{ti}\otimes\mathbf{x}_{ti}\Big)$ and $\zeta_{tb} \overset{\text{def}}{=} \frac{1}{b}\sum_{i=1}^{b}\xi_{ti} = \frac{1}{b}\sum_{i=1}^{b}\epsilon_{ti}\mathbf{x}_{ti}$. Equation 6 automatically brings out the "operator" view of analyzing the (expected) covariance of the centered estimate $\mathbb{E}\left[\eta_t\otimes\eta_t\right]$. We now note the following about the covariance of $\zeta_{tb}$:

$$\mathbb{E}\,\zeta_{tb}\otimes\zeta_{t'b} = \frac{1}{b^2}\sum_{i,j}\mathbb{E}\,\xi_{ti}\otimes\xi_{t'j}$$

$$= \Big[\frac{1}{b^2}\sum_{i=1}^{b}\mathbb{E}\,\xi_{ti}\otimes\xi_{ti}\Big]\mathbb{1}[t=t'] = \frac{1}{b}\mathbf{\Sigma}\;\;\mathbb{1}[t=t'] \tag{7}$$

Where, $\mathbb{1}[.]$ is the indicator function, and equals 1 if the argument inside $[.]$ is true and 0 otherwise. We note that the expectation of the cross terms in equation 7 is zero owing to independence of the samples $\{\mathbf{x}_{ti}, y_{ti}\}_{i=1}^{b}$ as well as between $\{\mathbf{x}_{ti}, y_{ti}\}_{i=1}^{b}$, $\{\mathbf{x}_{t'i}, y_{t'i}\}_{i=1}^{b}\;\forall\; t \neq t'$ and owing to the first order optimality conditions. Owing to the invariance of $\zeta_{tb}$ on the iteration $t$, context permitting, we sometimes drop the iteration index $t$ from $\zeta_{tb}$ and simply refer to it as $\zeta_b$.

Next we expand out the recurrence (6). Let $\mathbf{Q}_{j,t} = (\prod_{k=j}^{t}\mathbf{P}_{kb})^T$ with the convention that $\mathbf{Q}_{t',t} = \mathbf{I}\;\forall\; t' > t$. With this notation we have:

$$\begin{aligned}
\eta_t &= \mathbf{P}_{tb}\eta_{n-1} + \gamma\zeta_{tb} \\
&= \mathbf{P}_{tb}\mathbf{P}_{t-1,b}...\mathbf{P}_{1,b}\eta_0 + \gamma\sum_{j=0}^{t-1}\{\mathbf{P}_{tb}....\mathbf{P}_{t-j+1,b}\}\zeta_{t-j,b} \\
&= \mathbf{Q}_{1,t}\eta_0 + \gamma\sum_{j=0}^{t-1}\mathbf{Q}_{t-j+1,t}\zeta_{t-j,b} \\
&= \mathbf{Q}_{1,t}\eta_0 + \gamma\sum_{j=1}^{t}\mathbf{Q}_{j+1,t}\zeta_{j,b}\;.
\end{aligned} \tag{8}$$

### A.1.2 The terms in the averaged iterate

Let us consider beginning averaging after a certain number of iterations "$s$", i.e for $t > s$. In particular, after averaging the updates for the next "$N$" steps, we examine the quantity $\bar{\eta}_{s+1,N} = \overline{\mathbf{w}}_{s+1,N} - \mathbf{w}^*$, where $\overline{\mathbf{w}}_{s+1,N}$ is the iterate we obtain by averaging for "$N$" iterations starting from iteration $s+1$. We write out the expression for $\bar{\eta}_{s+1,N}$ starting out from equation 8:

$$\begin{aligned}
\bar{\eta}_{s+1,N} &= \frac{1}{N}\sum_{t=s+1}^{s+N}\eta_t \\
&= \frac{1}{N}\Big\{\sum_{t=s+1}^{s+N}\mathbf{Q}_{1,t}\Big\}\eta_0 + \frac{\gamma}{N}\sum_{t=s+1}^{s+N}\sum_{j=1}^{t}\mathbf{Q}_{j+1,t}\zeta_{j,b} \\
&= \frac{1}{N}\Big\{\sum_{t=s+1}^{s+N}\mathbf{Q}_{1,t}\Big\}\eta_0 + \frac{\gamma}{N}\Big\{\Big(\sum_{j=1}^{s}\sum_{t=s+1}^{s+N}\mathbf{Q}_{j+1,t}\zeta_{j,b}\Big) + \Big(\sum_{j=s+1}^{s+N}\sum_{t=j}^{s+N}\mathbf{Q}_{j+1,t}\zeta_{j,b}\Big)\Big\} \\
&= \frac{1}{N}\Big\{\sum_{t=s+1}^{s+N}\mathbf{Q}_{1,t}\Big\}\eta_0 + \frac{\gamma}{N}\sum_{j=1}^{s}\Big\{\sum_{t=s+1}^{s+N}\mathbf{Q}_{j+1,t}\Big\}\zeta_{j,b} + \frac{\gamma}{N}\sum_{j=s+1}^{s+N}\Big\{\sum_{t=j}^{s+N}\mathbf{Q}_{j+1,t}\Big\}\zeta_{j,b}
\end{aligned} \tag{9}$$

### A.1.3 Computing the outerproduct

Firstly, the outer product of $\bar{\eta}_{s+1,N}$ with itself contains three terms, each involving the outerproduct of each term in equation 9 with itself. The 6 other terms are the cross terms all of whose expectation are zero. We note that $\mathbb{E}[\xi_{t,i}] = \mathbb{E}[\epsilon_{t,i}\mathbf{x}_{t,i}] = 0\;\forall\; i,t$, and using this, it is clear to see the outer product of the cross term between terms 2 and 3 is zero, since they deal with disjoint samples $j = 1, ..., s$ and $j = s+1, ..., s+N$ and the $q_{.,.}$ doesn't deal with the same sample index $j$. The outer product of the $1^{st}$ term with the $2^{nd}$ and $3^{rd}$ term, $\mathbb{E}[x^{(i)}x^{(j)}x^{(k)}\epsilon] = 0$ (as noted by [9, 1], and if this is not satisfied we lose at most a constant factor two in our generalization error estimate), and as noted in section 4. This leaves us to deal with the 3 of the terms, involving outer product of the each of the terms with itself. Finally, we note that the bound in the theorem 5 refers to the outerproduct computation of the first term, and that of theorem 6 refers to the outerproduct computation of the second and third term.

### A.1.4 Proof of Theorem 5

We begin by noting that this proof generalizes the proof of Défossez and Bach [1] to account for mini-batching and tail-averaged SGD. In this proof, we assume that the problem is noiseless, i.e., we assume $\epsilon_{j,k} = 0$ a.s. $\forall\ j, k$.

$$\begin{aligned}
\mathbb{E}[\bar{\eta}_{s+1,N} \otimes \bar{\eta}_{s+1,N}] &= \frac{1}{N^2}\mathbb{E}\left[\Big\{\sum_{t=s+1}^{s+N}\mathbf{Q}_{1,t}\Big\}\eta_0 \otimes \eta_0\Big\{\sum_{t'=s+1}^{s+N}\mathbf{Q}_{1,t'}\Big\}^T\right] \\
&= \frac{1}{N^2}\mathbb{E}\left[\sum_{t,t'=s+1}^{s+N}\mathbf{Q}_{1,t}\eta_0 \otimes \eta_0\mathbf{Q}_{1,t'}^T\right] \\
&= \frac{1}{N^2}\sum_{t=s+1}^{s+N}\mathbb{E}\left[\mathbf{Q}_{1,t}\eta_0 \otimes \eta_0\mathbf{Q}_{1,t}^T + \sum_{t'=t+1}^{s+N}\Big\{\mathbf{Q}_{1,t}\eta_0 \otimes \eta_0\mathbf{Q}_{1,t}^T\mathbf{Q}_{t+1,t'}^T + \mathbf{Q}_{t+1,t'}\mathbf{Q}_{1,t}\eta_0 \otimes \eta_0\mathbf{Q}_{1,t}^T\Big\}\right] \\
&= \frac{1}{N^2}\sum_{t=s+1}^{s+N}\left[\mathbf{I} + \sum_{t'=t+1}^{s+N}(\mathbf{I}-\gamma\mathbf{H})_R^{(t'-t)} + (\mathbf{I}-\gamma\mathbf{H})_L^{(t'-t)}\right]\mathbb{E}[\mathbf{Q}_{1,t}\eta_0 \otimes \eta_0\mathbf{Q}_{1,t}^T] \quad (*) \\
&= \frac{1}{N^2}\sum_{t=s+1}^{s+N}\left[\mathbf{I} + \sum_{t'=t+1}^{s+N}(\mathbf{I}-\gamma\mathbf{H})_R^{(t'-t)} + (\mathbf{I}-\gamma\mathbf{H})_L^{(t'-t)}\right](\mathcal{I}-\gamma\mathcal{T}_b)^t\mathbf{E_0} \quad (**) \\
&= \frac{1}{N^2}\sum_{t=s+1}^{s+N}\Big[\mathbf{I} + (\gamma\mathcal{H}_\mathcal{R})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_R \\
&+ (\gamma\mathcal{H}_\mathcal{L})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_L\Big](\mathcal{I}-\gamma\mathcal{T}_b)^t\mathbf{E_0} \quad (***)
\end{aligned}$$

where, $(*)$ follows since $\mathbb{E}\left[Q_{t+1,t'}\right] = (\mathbf{I}-\gamma\mathbf{H})^{t'-t}$, and $(***)$ follows by summing up the geometric series (which converges when $\gamma \leq \frac{1}{\|\mathbf{H}\|_2}$, which is the case as discussed in section A.1.6), and finally, $(**)$ follows through the following argument:

$$\begin{aligned}
\mathbb{E}[\mathbf{Q}_{1,t}\eta_0 \otimes \eta_0\mathbf{Q}_{1,t}^T] &= \mathbb{E}[\mathbf{P}_{tb}\mathbf{P}_{t-1,b}...\mathbf{P}_{1b}\ \eta_0 \otimes \eta_0\ \mathbf{P}_{1b}...\mathbf{P}_{t-1,b}\mathbf{P}_{tb}] \\
&= \mathbb{E}[\mathbb{E}[\mathbf{P}_{tb}(\mathbf{P}_{t-1,b}...\mathbf{P}_{1b}\ \eta_0 \otimes \eta_0\ \mathbf{P}_{1b}...\mathbf{P}_{t-1,b})\mathbf{P}_{tb}|\mathcal{F}_{t-1}]]
\end{aligned}$$

where, $\mathcal{F}_{t-1}$ refers to the sigma-field defined by samples $\{\{(\mathbf{x}_{t',i}, y_{t',i})\}_{i=1}^b\}_{t'=1}^{t-1}$. We note that $(\mathbf{P}_{t-1,b}...\mathbf{P}_{1b}\ \eta_0 \otimes \eta_0\ \mathbf{P}_{1b}...\mathbf{P}_{t-1,b})$ is a constant w.r.t. the inner expectation because of the conditioning $\mathcal{F}_{t-1}$. Now, the observation can be seen clearly using the argument below (with $\mathbf{C}$ being a matrix that is independent w.r.t. the expectation):

$$\begin{aligned}
\mathbb{E}[\mathbf{P}_{tb}\mathbf{C}\mathbf{P}_{tb}^t] &= \mathbb{E}\Big[(\mathbf{I} - \frac{\gamma}{b}\sum_{i=1}^b\mathbf{x}_{ti}\otimes\mathbf{x}_{ti})\mathbf{C}(\mathbf{I} - \frac{\gamma}{b}\sum_{i=1}^b\mathbf{x}_{ti}\otimes\mathbf{x}_{ti})\Big] \\
&= \mathbf{C} - \gamma\mathbf{C}\mathbf{H} - \gamma\mathbf{H}\mathbf{C} + \frac{\gamma^2}{b}E[\mathbf{x}^T\mathbf{C}\mathbf{x}\ \mathbf{x}\otimes\mathbf{x}] + \frac{\gamma^2(b-1)}{b}E[\mathbf{x}_1^T\mathbf{C}\mathbf{x}_2\ \mathbf{x}_1\otimes\mathbf{x}_2] \\
&= (\mathcal{I} - \gamma(H_L + H_R - \frac{\gamma}{b}\mathcal{M} - \frac{\gamma(b-1)}{b}\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R}))\mathbf{C} \\
&= (\mathcal{I} - \gamma\mathcal{T}_b)\mathbf{C} \qquad (10)
\end{aligned}$$

where, we apply the fact that the samples $\{(\mathbf{x}_{ti}, y_{ti})\}_{i=1}^b$ are independent and identically distributed draws from the distribution $\mathcal{D}$. The final expression is obtained by applying the argument recursively for all $t' = t, t-1, .., 1$.

Next we bound the exponential term, in particular:

$$\begin{aligned}
\mathbf{A}_N &= \frac{-1}{N^2}\sum_{t=s+1}^{s+N}\Big[(\gamma\mathcal{H}_\mathcal{R})^{-1}(\mathbf{I}-\gamma\mathbf{H})_R^{s+N-t+1} + (\gamma\mathcal{H}_\mathcal{L})^{-1}(\mathbf{I}-\gamma\mathbf{H})_L^{s+N-t+1}\Big](\mathcal{I}-\gamma\mathcal{T}_b)^t\mathbf{E_0} \\
\implies \|\mathbf{A}_N\|_F &\leq \frac{2}{\gamma\lambda_{\min}(\mathbf{H})N}\rho_b^{s+N+1}\|\mathbf{E_0}\|_F
\end{aligned}$$

The remaining terms are bounded as the following:

$$\begin{aligned}\mathbb{E}[\bar{\eta}_{s+1,N} \otimes \bar{\eta}_{s+1,N}] &= \frac{1}{\gamma N^2} \sum_{t=s+1}^{s+N} [{\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma \mathcal{I}](\mathcal{I} - \gamma \mathcal{T}_b)^t \mathbf{E_0} + \mathbf{A}_N \\ &= \frac{1}{\gamma^2 N^2} [{\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma \mathcal{I}] \mathcal{T}_b^{-1} [(\mathcal{I} - \gamma \mathcal{T}_b)^{s+1} - (\mathcal{I} - \gamma \mathcal{T}_b)^{s+N+1}] \mathbf{E_0} + \mathbf{A}_N\end{aligned}$$

Considering the exponential terms separately, we write out the covariance operator as:

$$\mathbb{E}[\bar{\eta}_{s+1,N} \otimes \bar{\eta}_{s+1,N}] = \frac{1}{\gamma^2 N^2} [{\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma \mathcal{I}] \mathcal{T}_b^{-1} (\mathcal{I} - \gamma \mathcal{T}_b)^{s+1} \mathbf{E_0} + \mathbf{A}_N + \mathbf{B}_N$$

with $\mathbf{B}_N$ defined as:

$$\begin{aligned}\mathbf{B}_N &= \frac{-1}{\gamma^2 N^2} [{\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma \mathcal{I}] \mathcal{T}_b^{-1} (\mathcal{I} - \gamma \mathcal{T}_b)^{s+N+1} \mathbf{E_0} \\ \implies \|\mathbf{B}_N\|_F &\leq \frac{1}{\gamma^2 N^2 \lambda_{\min}(\mathcal{T}_b)} \Big( \frac{2}{\lambda_{\min}(\mathbf{H})} - \gamma \Big) \rho_{Tb}^{s+N+1} \|\mathbf{E_0}\|_F\end{aligned}$$

The last equation above assumes that $\gamma < \frac{2}{\lambda_{\min}(\mathbf{H})}$. Finally, using triangle inequality that $\|\mathbf{A}_N + \mathbf{B}_N\|_F \leq \|\mathbf{A}_N\|_F + \|\mathbf{B}_N\|_F$ and hence, when we begin averaging starting at iteration $s+1$ and average for $N$ iterations subsequently, we get a geometric decay of the bias term. The final covariance operator of the bias term thus has the following asymptotic expansion:

$$\mathbb{E}[\bar{\eta}_{s+1,N} \otimes \bar{\eta}_{s+1,N}] = \frac{1}{\gamma^2 N^2} [{\mathcal{H}_{\mathcal{L}}}^{-1} + {\mathcal{H}_{\mathcal{R}}}^{-1} - \gamma \mathcal{I}] \mathcal{T}_b^{-1} (\mathcal{I} - \gamma \mathcal{T}_b)^{s+1} \mathbf{E_0} + \mathcal{O}(\frac{\rho_b^{s+N+1}}{\gamma \lambda_{\min}(\mathbf{H}) N} \|\mathbf{E_0}\|_F)$$

The above bound requires the following conditions to hold: (i) $\mathcal{T}_b(\gamma) \succeq 0$, (ii) $\gamma < \frac{1}{\lambda_{\max}(\mathbf{H})}$. This wraps up the proof of Theorem 5 and indicates the geometric decay of the bias term during the unaveraged phase and the $\mathcal{O}(\frac{1}{N^2})$ decay during the averaging phase.

### A.1.5 Proof of Theorem 6

We begin by noting that this proof generalizes the proof of Défossez and Bach [1] to account for mini-batching and tail-averaged SGD. For bounding the variance part of the generalization error, we go over a sharp analysis of the stochastic process defined by the mini-batch tail averaged SGD initialized at the solution $\mathbf{w}^*$, i.e. $\mathbf{w}_0 = \mathbf{w}^*$, and thus the process is driven solely by the noise. We note that we require bounding the behavior of two terms, the first of which is defined by the noise introduced during the unaveraged phase, and the second being one that deals with the noise introduced during the tail averaging phase.

$$\begin{aligned}\mathbb{E}[\bar{\eta}_{s+1,N} \otimes \bar{\eta}_{s+1,N}] &= \frac{\gamma^2}{N^2} \left( \sum_{j=1}^{s} \sum_{t,t'=s+1}^{s+N} \mathbb{E}[\mathbf{Q}_{j+1,t} \zeta_{jb} \otimes \zeta_{jb} \mathbf{Q}_{j+1,t'}^T] + \sum_{j=s+1}^{s+N} \sum_{t,t'=j}^{s+N} \mathbb{E}[\mathbf{Q}_{j+1,t} \zeta_{jb} \otimes \zeta_{jb} \mathbf{Q}_{j+1,t'}^T] \right) \\ &= \frac{\gamma^2}{N^2} (\mathbf{K}_1 + \mathbf{K}_2) \qquad (11)\end{aligned}$$

We consider the first term, and specifically the inner two terms of the summation:

$$\begin{aligned}\mathbf{K}_{1j} &= \sum_{t=s+1}^{s+N} \sum_{t'=s+1}^{s+N} \mathbb{E}[\mathbf{Q}_{j+1,t} \zeta_{jb} \otimes \zeta_{jb} \mathbf{Q}_{j+1,t'}^T] \\ &= \sum_{t=s+1}^{s+N} \mathbb{E}\Big[ \mathbf{Q}_{j+1,t} \zeta_{jb} \otimes \zeta_{jb} \mathbf{Q}_{j+1,t}^T + \sum_{t'=t+1}^{s+N} \mathbf{Q}_{j+1,t} \zeta_{jb} \otimes \zeta_{jb} \mathbf{Q}_{j+1,t}^T \mathbf{Q}_{t+1,t'}^T + \sum_{t'=s+1}^{t-1} \mathbf{Q}_{t'+1,t} \mathbf{Q}_{j+1,t'} \zeta_{jb} \otimes \zeta_{jb} \mathbf{Q}_{j+1,t'}^T \Big]\end{aligned}$$

$$= \sum_{t=s+1}^{s+N} \mathbb{E}\Big[\mathbf{Q}_{j+1,t}\zeta_{jb} \otimes \zeta_{jb}\mathbf{Q}_{j+1,t}^T + \sum_{t'=t+1}^{s+N} \mathbf{Q}_{j+1,t}\zeta_{jb} \otimes \zeta_{jb}\mathbf{Q}_{j+1,t}^T(\mathbf{I}-\gamma\mathbf{H})^{t'-t} + \sum_{t'=s+1}^{t-1} (\mathbf{I}-\gamma\mathbf{H})^{t-t'}\mathbf{Q}_{j+1,t'}\zeta_{jb} \otimes \zeta_{jb}\mathbf{Q}_{j+1,t'}^T\Big]$$

$$= \sum_{t=s+1}^{s+N} \Big[\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_L\Big](\mathcal{I}-\gamma\mathcal{T}_b)^{t-j}\mathbb{E}[\zeta_{jb} \otimes \zeta_{jb}]$$

$$= \frac{1}{b}\sum_{t=s+1}^{s+N} \Big[\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_L\Big](\mathcal{I}-\gamma\mathcal{T}_b)^{t-j}\mathbf{\Sigma}$$

Note, the second to last line above is obtained by seeing that we can swap the summations for the third term and then replace the roles of $t$ and $t'$ and applying equation 10, following which we note that $\mathbb{E}[\zeta_{jb} \otimes \zeta_{jb}] = \frac{1}{b}\mathbf{\Sigma}$ from equation 7 for the equation in the last line above. And hence, the first term $\mathbf{K}_1$ is:

$$\mathbf{K}_1 = \sum_{j=1}^{s} \mathbf{K}_{1j}$$

$$= \frac{1}{b}\sum_{j=1}^{s}\sum_{t=s+1}^{s+N} \Big[\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_L\Big](\mathcal{I}-\gamma\mathcal{T}_b)^{t-j}\mathbf{\Sigma} \tag{12}$$

In a very similar manner, it is easy to observe the second term is:

$$\mathbf{K}_2 = \frac{1}{b}\sum_{j=s+1}^{s+N}\sum_{t=j}^{s+N} \Big[\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_L\Big](\mathcal{I}-\gamma\mathcal{T}_b)^{t-j}\mathbf{\Sigma} \tag{13}$$

If $\mathbf{K}_1$ and $\mathbf{K}_2$ can be bounded, then it is clear that the variance decays as $\frac{\gamma^2}{N^2}(\mathbf{K}_1 + \mathbf{K}_2)$ as is seen from equation 11. Let us begin with $\mathbf{K}_1$ (from equation 12). We first exchange the two summations and sum up the resulting series w.r.t. $j$.

$$\mathbf{K}_1 = \frac{1}{b}\sum_{t=s+1}^{s+N} \Big[\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}\{(\mathbf{I}-\gamma\mathbf{H}) -$$
$$(\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}\}_L\Big](\gamma\mathcal{T}_b)^{-1}\Big[(\mathcal{I}-\gamma\mathcal{T}_b)^{t-s} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t}\Big]\mathbf{\Sigma}$$
$$= \frac{1}{b}\sum_{t=1}^{N} \Big[\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{N-t+1}]_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) -$$
$$(\mathbf{I}-\gamma\mathbf{H})^{N-t+1}]_L\Big](\gamma\mathcal{T}_b)^{-1}\Big[(\mathcal{I}-\gamma\mathcal{T}_b)^{t} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t+s}\Big]\mathbf{\Sigma}$$

Next we bound the exponential terms, i.e.

$$\mathbf{K}_1 = \frac{1}{b}\sum_{t=1}^{N} \Big[(\gamma\mathcal{H}_{\mathcal{L}})^{-1} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-1}\Big[(\mathcal{I}-\gamma\mathcal{T}_b)^{t} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t+s}\Big]\mathbf{\Sigma} + \mathbf{A}'_N$$
$$= \frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{L}})^{-1} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}\Big[(\mathcal{I}-\gamma\mathcal{T}_b) - (\mathcal{I}-\gamma\mathcal{T}_b)^{s+1}\Big]\Big[\mathcal{I} - (\mathcal{I}-\gamma\mathcal{T}_b)^{N}\Big]\mathbf{\Sigma} + \mathbf{A}'_N$$

where $\mathbf{A}'_N$ is:

$$\mathbf{A}'_N = -\frac{1}{b}\sum_{t=1}^{N} \Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_R^{N-t+1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_L^{N-t+1}\Big](\gamma\mathcal{T}_b)^{-1}\Big[(\mathcal{I}-\gamma\mathcal{T}_b)^{t} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t+s}\Big]\mathbf{\Sigma}$$

whose norm is bounded as:

$$\|\mathbf{A}'_N\|_F \leq \frac{2N}{\gamma^2 b\lambda_{\min}(\mathbf{H})\lambda_{\min}(\mathcal{T}_b)}(\rho_b^{N+1} + \rho_b^{N+s+1})\|\mathbf{\Sigma}\|_F$$

Again, we accumulate the exponential terms existing in $\mathbf{K}_1$ in $B'_N$, leading to:

$$\begin{aligned}\mathbf{K}_1 &= \frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{L}})^{-1} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}\Big[(\mathcal{I}-\gamma\mathcal{T}_b) - (\mathcal{I}-\gamma\mathcal{T}_b)^{s+1}\Big]\mathbf{\Sigma} + \mathbf{B}'_N + \mathbf{A}'_N \\ &= \frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{L}})^{-1} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}\Big[(\mathcal{I}-\gamma\mathcal{T}_b)\Big]\mathbf{\Sigma} + \mathbf{C}'_N + \mathbf{B}'_N + \mathbf{A}'_N \end{aligned} \tag{14}$$

where, $\mathbf{B}'_N$ is:

$$\mathbf{B}'_N = -\frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{L}})^{-1} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}\Big[(\mathcal{I}-\gamma\mathcal{T}_b) - (\mathcal{I}-\gamma\mathcal{T}_b)^{s+1}\Big](\mathcal{I}-\gamma\mathcal{T}_b)^{N}\mathbf{\Sigma}$$

whose norm is bounded using triangle inequality as:

$$\|\mathbf{B}'_N\|_F \leq \frac{1}{b\gamma^3(\lambda_{\min}(\mathcal{T}_b))^2}\Big(\frac{2}{\lambda_{\min}(\mathbf{H})} - \gamma\Big)(\rho_{Tb} + \rho_{Tb}^{s+1})\rho_{Tb}^{N}\|\mathbf{\Sigma}\|_F$$

finally, $\mathbf{C}'_N$ is:

$$\mathbf{C}'_N = -\frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{L}})^{-1} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}(\mathcal{I}-\gamma\mathcal{T}_b)^{s+1}\mathbf{\Sigma}$$

whose norm is bounded as:

$$\|\mathbf{C}'_N\| \leq \frac{1}{b\gamma^3(\lambda_{\min}(\mathcal{T}_b))^2}\Big(\frac{2}{\lambda_{\min}(\mathbf{H})} - \gamma\Big)\rho_{Tb}^{s+1}\|\mathbf{\Sigma}\|_F$$

which rounds up bounding all terms of equation 14, except for the leading order term. This term ends up canceling with a term involved in bounding $\mathbf{K}_2$ in equation 17.

We now begin bounding the terms in $\mathbf{K}_2$ (from equation 13). Towards this effect, we begin by interchanging the summations in $\mathbf{K}_2$:

$$\begin{aligned}\mathbf{K}_2 = \frac{1}{b}\sum_{t=s+1}^{s+N}\sum_{j=s+1}^{t}\Big[&\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}]_R + \\ &(\gamma\mathcal{H}_{\mathcal{L}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}]_L\Big](\mathcal{I}-\gamma\mathcal{T}_b)^{t-j}\mathbf{\Sigma} \\ = \frac{1}{b}\sum_{t=s+1}^{s+N}\Big[&\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}]_R + \\ &(\gamma\mathcal{H}_{\mathcal{L}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{s+N-t+1}]_L\Big](\gamma\mathcal{T}_b)^{-1}[\mathcal{I} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t-s}]\mathbf{\Sigma} \\ = \frac{1}{b}\sum_{t=1}^{N}\Big[&\mathcal{I} + (\gamma\mathcal{H}_{\mathcal{R}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{N-t+1}]_R + \\ &(\gamma\mathcal{H}_{\mathcal{L}})^{-1}[(\mathbf{I}-\gamma\mathbf{H}) - (\mathbf{I}-\gamma\mathbf{H})^{N-t+1}]_L\Big](\gamma\mathcal{T}_b)^{-1}\Big[\mathcal{I} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t}\Big]\mathbf{\Sigma}\end{aligned}$$

We begin considering the exponential terms, i.e.,:

$$\mathbf{K}_2 = \frac{1}{b}\sum_{t=1}^{N}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-1}\Big[\mathcal{I} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t}\Big]\mathbf{\Sigma} + \mathbf{A}''_N$$

where $\mathbf{A}''_N$ is:

$$\begin{aligned}\mathbf{A}''_N &= -\frac{1}{b}\sum_{t=1}^{N}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_R^{N-t+1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_L^{N-t+1}\Big](\gamma\mathcal{T}_b)^{-1}\Big[\mathcal{I} - (\mathcal{I}-\gamma\mathcal{T}_b)^{t}\Big]\mathbf{\Sigma} \\ &= \mathbf{B}''_N + \mathbf{C}''_N\end{aligned}$$

where,

$$\begin{aligned}\mathbf{B}''_N &= -\frac{1}{b}\sum_{t=1}^{N}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_R^{N-t+1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_L^{N-t+1}\Big](\gamma\mathcal{T}_b)^{-1}\mathbf{\Sigma} \\ &= -\frac{1}{b}\left[(\gamma\mathcal{H}_{\mathcal{R}})^{-2}\Big[(\mathbf{I}-\gamma\mathbf{H})_R - (\mathbf{I}-\gamma\mathbf{H})_R^{N+1}\Big] + (\gamma\mathcal{H}_{\mathcal{L}})^{-2}\Big[(\mathbf{I}-\gamma\mathbf{H})_L - (\mathbf{I}-\gamma\mathbf{H})_L^{N+1}\Big]\right](\gamma\mathcal{T}_b)^{-1}\mathbf{\Sigma} \\ &= \mathbf{F}''_N + \mathbf{G}''_N\end{aligned}$$

Where $\mathbf{F}''_N$ is:

$$\begin{aligned}\mathbf{F}''_N &= -\frac{1}{b}\left[(\gamma\mathcal{H}_{\mathcal{R}})^{-2}(\mathbf{I}-\gamma\mathbf{H})_R + (\gamma\mathcal{H}_{\mathcal{L}})^{-2}(\mathbf{I}-\gamma\mathbf{H})_L\right](\gamma\mathcal{T}_b)^{-1}\mathbf{\Sigma} \\ &= -\frac{1}{b\gamma^3}\left[(\mathcal{H}_{\mathcal{R}}{}^{-2} - \gamma\mathcal{H}_{\mathcal{R}}{}^{-1}) + (\mathcal{H}_{\mathcal{L}}{}^{-2} - \gamma\mathcal{H}_{\mathcal{L}}{}^{-1})\right]\mathcal{T}_b{}^{-1}\mathbf{\Sigma} \\ &= -\frac{1}{b\gamma^3}\left[(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1} - \gamma\mathcal{I})(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1}) - \mathcal{H}_{\mathcal{L}}{}^{-1}\mathcal{H}_{\mathcal{R}}{}^{-1} - \mathcal{H}_{\mathcal{R}}{}^{-1}\mathcal{H}_{\mathcal{L}}{}^{-1}\right]\mathcal{T}_b{}^{-1}\mathbf{\Sigma} \\ &= -\frac{1}{b\gamma^3}\left[(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1} - \gamma\mathcal{I})(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1}) - 2\mathcal{H}_{\mathcal{L}}{}^{-1}\mathcal{H}_{\mathcal{R}}{}^{-1}\right]\mathcal{T}_b{}^{-1}\mathbf{\Sigma} \qquad (15)\end{aligned}$$

The norm of $\mathbf{F}''_N$ is bounded along with the leading order variance term. $\mathbf{G}''_N$ is then written out as:

$$\mathbf{G}''_N = \frac{1}{b}\left[(\gamma\mathcal{H}_{\mathcal{R}})^{-2}(\mathbf{I}-\gamma\mathbf{H})_R^{N+1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-2}(\mathbf{I}-\gamma\mathbf{H})_L^{N+1}\right](\gamma\mathcal{T}_b)^{-1}\mathbf{\Sigma}$$

whose norm is bounded as:

$$\|\mathbf{G}''_N\|_F \leq \frac{\rho_{Hb}^{N+1}}{b\gamma^3(\lambda_{\min}(\mathbf{H}))^2\lambda_{\min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F$$

$$\mathbf{C}''_N = \frac{1}{b}\sum_{t=1}^{N}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_R^{N-t+1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1}(\mathbf{I}-\gamma\mathbf{H})_L^{N-t+1}\Big](\gamma\mathcal{T}_b)^{-1}(\mathcal{I}-\gamma\mathcal{T}_b)^t\mathbf{\Sigma}$$

whose norm is bounded as:

$$\|\mathbf{C}''_N\|_F \leq \frac{2}{\gamma^2 b\lambda_{\min}(\mathbf{H})\lambda_{\min}(\mathcal{T}_b)}(N\rho_b^{N+1})\|\mathbf{\Sigma}\|_F$$

Now, $\mathbf{K}_2$ is written as the following:

$$\mathbf{K}_2 = \frac{N}{b}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-1}\mathbf{\Sigma} + \mathbf{D}''_N + \mathbf{A}''_N \qquad (16)$$

where the $N$ upfront is the time for which the iterates have been averaged, and this is the term that leads to $\frac{1}{N}$ rate of decay for the variance term (which has the factor of $\frac{\gamma^2}{N^2}\mathbf{K}_2$), and $\mathbf{D}''_N$ is obtained through summing the resulting geometric series:

$$\begin{aligned}\mathbf{D}''_N &= -\frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}\Big[(\mathcal{I}-\gamma\mathcal{T}_b) - (\mathcal{I}-\gamma\mathcal{T}_b)^{N+1}\Big]\mathbf{\Sigma} \\ &= -\frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}(\mathcal{I}-\gamma\mathcal{T}_b)\mathbf{\Sigma} + \mathbf{E}''_N \qquad (17)\end{aligned}$$

where, $\mathbf{E}''_N$ is:

$$\mathbf{E}''_N = \frac{1}{b}\Big[(\gamma\mathcal{H}_{\mathcal{R}})^{-1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1} - \mathcal{I}\Big](\gamma\mathcal{T}_b)^{-2}(\mathcal{I}-\gamma\mathcal{T}_b)^{N+1}\mathbf{\Sigma}$$

whose norm is bounded as:

$$\|\mathbf{E}''_N\|_F \leq \frac{1}{\gamma^3 b(\lambda_{\min}(\mathcal{T}_b))^2}\Big(\frac{2}{\lambda_{\min}(\mathbf{H})} - \gamma\Big)\rho_{Tb}^{N+1}\|\mathbf{\Sigma}\|_F$$

We note that the leading term in $\mathbf{D}_N''$ in equation 17 is exactly the negative of the leading order term in equation 14, thus canceling out without requiring us to bound their behavior. This allows us to provide the final theorem statement, combining the leading order terms from equations 15 and 16 along with the $\frac{\gamma^2}{N^2}$ outside the summation,

$$\begin{aligned}
\mathbb{E}[\bar{\eta}_{s,N} \otimes \bar{\eta}_{s,N}] &= \frac{1}{bN}\left[(\gamma\mathcal{H}_{\mathcal{R}})^{-1} + (\gamma\mathcal{H}_{\mathcal{L}})^{-1} - \mathcal{I}\right](\gamma\mathcal{T}_b)^{-1}\mathbf{\Sigma} \\
&- \frac{1}{b\gamma N^2}\left[(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1} - \gamma\mathcal{I})(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1}) - 2\mathcal{H}_{\mathcal{L}}{}^{-1}\mathcal{H}_{\mathcal{R}}{}^{-1}\right]\mathcal{T}_b{}^{-1}\mathbf{\Sigma} \\
&+ \mathcal{O}(\frac{\rho_b^{N+1}}{bN\lambda_{\min}(\mathbf{H})\lambda_{\min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F) \\
&= \frac{1}{bN}\left[\mathcal{I} - \frac{1}{\gamma N}(\mathcal{H}_{\mathcal{L}}{}^{-1} + \mathcal{H}_{\mathcal{R}}{}^{-1})\right]\left[(\mathcal{H}_{\mathcal{R}})^{-1} + (\mathcal{H}_{\mathcal{L}})^{-1} - \gamma\mathcal{I}\right]\mathcal{T}_b{}^{-1}\mathbf{\Sigma} \\
&+ \frac{2}{\gamma bN^2}\mathcal{H}_{\mathcal{L}}{}^{-1}\mathcal{H}_{\mathcal{R}}{}^{-1}\mathcal{T}_b{}^{-1}\mathbf{\Sigma} + \mathcal{O}(\frac{\rho_b^{N+1}}{bN\lambda_{\min}(\mathbf{H})\lambda_{\min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F)
\end{aligned}$$

with which the proof of Theorem 6 concludes.

### A.1.6 Some intermediate Lemmas required to prove Lemma 7

We establish a sequence of lemmas in this subsection in order to prove that the operators $\mathbf{I} - \gamma\mathbf{H}$ and $\mathcal{I} - \gamma\mathcal{T}_b$ contract with the choice of our stepsizes as described in lemma 7. We then provide a scalar bound for the contraction factor, and this is done by generalizing the lemma appearing in Défossez and Bach [1] to the mini-batching case.

**Lemma 8.** *Without loss of generality, let us assume that the eigenvalues $\{\lambda_i\}_{i=1}^d > 0$ of $\mathbf{H}$ be arranged in decreasing order of their values. If $\gamma > 0$ is chosen such that $\mathcal{T}_b(\gamma) \succ 0$, i.e.:*

$$\langle \mathbf{A}, \mathcal{T}_b\mathbf{A}\rangle = 2\,\mathrm{Tr}\,\mathbf{A}^T\mathbf{H}\mathbf{A} - \frac{\gamma}{b}\,\mathrm{Tr}\,\mathbf{A}^T\mathcal{M}\mathbf{A} - \gamma\frac{b-1}{b}\,\mathrm{Tr}\,\mathbf{A}^T\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A} > 0\ \forall\ \mathbf{A} \in\ \mathcal{S}(d)$$

*then the following upper bound on the stepsize $\gamma$ holds: $\gamma < \frac{b}{b+1}(\frac{\lambda_i+\lambda_j}{\lambda_i\lambda_j})$ holds $\forall\ \{i,j\} \in \{1,2,...,d\}\times\{1,2,...,d\}$*

*Proof.* This proof of this lemma generalizes the proof appearing in Défossez and Bach [1] to establish bounds on the stepsize for the mini-batching case.

Expanding out the inner product, we have:

$$\begin{aligned}
\langle \mathbf{A}, \mathcal{T}_b\mathbf{A}\rangle &= \mathrm{Tr}(\mathbf{A}^T(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma}{b}\mathcal{M} - \gamma(\frac{b-1}{b})\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})\mathbf{A}) \\
&= 2\,\mathrm{Tr}(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}\,\mathrm{Tr}(\mathbf{A}^T\mathcal{M}\mathbf{A}) - \gamma(\frac{b-1}{b})\,\mathrm{Tr}(\mathbf{A}^T\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A}) \\
&= 2\,\mathrm{Tr}(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}\mathbb{E}(\mathbf{x}^T\mathbf{A}\mathbf{x})^2 - \gamma(\frac{b-1}{b})\mathbb{E}(\mathbf{x}_1^T\mathbf{A}\mathbf{x}_2)^2
\end{aligned}$$

Where, $(\mathbf{x}_1, y_1)$, $(\mathbf{x}_2, y_2)$ and $(\mathbf{x}, y)$ are independent draws from the distribution $\mathcal{D}$. Next, we consider the following:

$$\begin{aligned}
2\,\mathrm{Tr}\,(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}[\mathrm{Tr}\,(\mathbf{A}\mathbf{H})]^2 - \gamma(\frac{b-1}{b})\,\mathrm{Tr}(\mathbf{A}^T\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A}) &= 2\,\mathrm{Tr}\,(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}[\mathrm{Tr}\,\mathbf{A}\mathbb{E}\,[\mathbf{x}\otimes\mathbf{x}]]^2 - \gamma(\frac{b-1}{b})\,\mathrm{Tr}(\mathbf{A}\mathbf{H}\mathbf{A}\mathbf{H}) \\
&\geq 2\,\mathrm{Tr}\,(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}\mathbb{E}(\mathbf{x}^T\mathbf{A}\mathbf{x})^2 - \gamma(\frac{b-1}{b})\,\mathrm{Tr}(\mathbf{A}^T\mathbf{H}\mathbf{A}\mathbf{H}) \\
&= \langle \mathbf{A}, \mathcal{T}_b\mathbf{A}\rangle > 0
\end{aligned}$$

Where, we note that the second line follows through the application of Jensen's inequality. Further, Let $\mathbf{H} = \mathbf{U}\mathbf{D}\mathbf{U}^t$ to be the eigen decomposition of $\mathbf{H}$ so that $\mathbf{D}$ contains the eigenvalues $\{\lambda_i\}_{i=1}^d$ in its diagonal. For any $(i,j) \in \{1,...,d\}\times\{1,...,d\}$, we set the matrix $\mathbf{A} = c_i\mathbf{u}_i \otimes \mathbf{u}_i + c_j\mathbf{u}_j \otimes \mathbf{u}_j$ so as to observe the following:

$$\begin{aligned}
2\,\mathrm{Tr}\,(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}[\mathrm{Tr}\,\mathbf{A}\mathbf{H}]^2 - \gamma(\frac{b-1}{b})\,\mathrm{Tr}(\mathbf{A}^T\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A}) &= 2(c_i^2\lambda_i + c_j^2\lambda_j) - \frac{\gamma}{b}(c_i\lambda_i + c_j\lambda_j)^2 - \gamma(\frac{b-1}{b})(c_i^2\lambda_i^2 + c_j^2\lambda_j^2) \\
&> 0
\end{aligned}$$

By setting $c_i = \sqrt{\frac{\lambda_j}{\lambda_i}}$ and $c_j = \sqrt{\frac{\lambda_i}{\lambda_j}}$, the result follows. □

**Corollary 8.1.** *By setting $i = j = 1$ in lemma 8, we observe that the condition $\mathcal{T}_b \succ 0$ allows $\gamma < (\frac{b}{b+1}) \cdot \frac{2}{\lambda_{max}(\mathbf{H})}$.*

*Remark* 2. We note that an immediate consequence of corollary 8.1 is that by setting $\gamma$ such that $\mathcal{T}_b \succ 0$, $-\mathbf{I} \prec \mathbf{I} - \gamma\mathbf{H} \prec \mathbf{I}$.

**Lemma 9.** *With $\mathcal{T}_b = \mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma}{b}\mathcal{M} - \gamma(\frac{b-1}{b})\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}$ defined as before, if $\gamma > 0$ is chosen such that $\langle \mathbf{Z}, \mathcal{T}_b\mathbf{Z}\rangle > 0 \; \forall \; \mathbf{Z} \in \mathcal{S}(d)$, then, $\forall \; d \geq 2$, we have:*

$$\langle \mathbf{Z}, (\mathcal{I} - \gamma\mathcal{T}_b)\,\mathbf{Z}\rangle \geq -\langle \mathbf{Z}, \mathbf{Z}\rangle \; \forall \; \mathbf{Z} \in \mathcal{S}(d).$$

*while in dimension $d = 1$ we have the following condition:*

$$\mathcal{I} - \gamma\mathcal{T}_b \geq 0$$

Before we prove the above lemma, we prove a simpler version of this lemma where a stronger conclusion holds, but only $\forall \; \mathbf{Z} \in \mathcal{S}^+(\mathbb{R}^d)$ (as opposed to the statement of lemma 9 where $\mathbf{Z} \in \mathcal{S}(d)$).

**Lemma 10.** *If $\gamma < \frac{1}{\|\mathbf{H}\|_2}$, the following statement holds:*

$$\langle \mathbf{Z}, (\mathcal{I} - \gamma(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \gamma\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}))\,\mathbf{Z}\rangle \geq 0 \; \forall \; \mathbf{Z} \in \mathcal{S}(d).$$

*Proof.* Assuming the eigenvalues of $\mathbf{H}$ are $\{\lambda_i\}_{i=1}^d$ arranged in decreasing order, the proof requires the observation that the eigenvalues of $(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \gamma b\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})$ are respectively $\{\lambda_i + \lambda_j - \gamma\lambda_i\lambda_j\}_{i,j=1}^d$. This is because if $\{\mathbf{u}_i\}_{i=1}^d$ are the eigenvectors of $\mathbf{H}$, then, $\{\mathbf{u}_i \otimes \mathbf{u}_j\}_{i,j=1}^d$ are the eigenvectors for $\mathcal{H}_{\mathcal{L}}$, $\mathcal{H}_{\mathcal{R}}$ and $\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}$ with corresponding eigenvalues being $\lambda_i$, $\lambda_j$ and $\lambda_i\lambda_j$ respectively. This is because:

$$\begin{aligned}
\mathcal{H}_{\mathcal{L}}\mathbf{u}_i \otimes \mathbf{u}_j &= \mathbf{H}\mathbf{u}_i \otimes \mathbf{u}_j = \lambda_i\mathbf{u}_i \otimes \mathbf{u}_j \\
\mathcal{H}_{\mathcal{R}}\mathbf{u}_i \otimes \mathbf{u}_j &= \mathbf{u}_i \otimes \mathbf{u}_j\mathbf{H} = \lambda_j\mathbf{u}_i \otimes \mathbf{u}_j \\
\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{u}_i \otimes \mathbf{u}_j &= \mathbf{H}\mathbf{u}_i \otimes \mathbf{u}_j\mathbf{H} = \lambda_i\lambda_j\mathbf{u}_i \otimes \mathbf{u}_j
\end{aligned}$$

If we prove that the eigenvalues of $(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \gamma\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})$ are all upper bounded by $\frac{1}{\gamma}$, the lemma is proven since this implies $\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \gamma\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}} \prec \frac{1}{\gamma}\mathcal{I}$.

We consider two cases, (i) when $i = j$, in which case we deal with eigenvalues of the form $2\lambda_i - \gamma\lambda_i^2$. If we require this to be less than $\frac{1}{\gamma}$, then, $(\gamma\lambda_i - 1)^2 > 0$ which holds automatically $\forall\; \gamma$. (ii) When $i \neq \; j$, we require $\lambda_i + \lambda_j - \gamma\lambda_i\lambda_j < \frac{1}{\gamma}$, or, $\gamma^2\lambda_i\lambda_j - \gamma(\lambda_i + \lambda_j) + 1 > 0$. If we assume without loss of generality, $\lambda_i > \lambda_j$, then, the solution of interest is $\gamma < \frac{1}{\lambda_i}$. Taking the intersection of all the solutions obtained from solving equations $\forall\; i, j$, we obtain, $\gamma < \frac{1}{\lambda_{\max}(\mathbf{H})}$. We consider lemma 8 with $i = 1$ and $j \neq \; 1$. In this case, we know that $\gamma < \frac{b(\lambda_{\max}(\mathbf{H}) + \lambda_j)}{(b+1)\lambda_{\max}(\mathbf{H})\lambda_j}$. The proof completes iff $\frac{1}{\lambda_{\max}(\mathbf{H})} < \frac{b(\lambda_{\max}(\mathbf{H}) + \lambda_j)}{(b+1)\lambda_{\max}(\mathbf{H})\lambda_j}$ or $\frac{b+1}{b} < \frac{\lambda_j + \lambda_{\max}(\mathbf{H})}{\lambda_j}$ implying $b > \frac{\lambda_j}{\lambda_{\max}(\mathbf{H})}$, which is true $\forall\; b \geq 1$. We also note that the result holds even in the case that there exists eigenvalues with multiplicity $\geq \; 1$ since that is covered through case(i). □

**Lemma 11.** *For any PSD matrix $\mathbf{A} \in \mathcal{S}^+(\mathbb{R}^d)$, the following inequality holds:*

$$\langle \mathbf{A}, \mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A}\rangle \leq \langle \mathbf{A}, \mathcal{M}\mathbf{A}\rangle$$

*Proof.* The claim is proven through the following argument:

$$\begin{aligned}
\langle \mathbf{A}, \mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A}\rangle &= \langle \mathbf{A}, \mathbf{H}\mathbf{A}\mathbf{H}\rangle \\
&= \mathbb{E}_{\mathbf{x}_1, \mathbf{x}_2}\left[(\mathbf{x}_1{}^\top\mathbf{A}\mathbf{x}_2)^2\right] \\
&\leq \left(\mathbb{E}_{\mathbf{x}}\left[\mathbf{x}^\top\mathbf{A}\mathbf{x}\right]\right)^2 \\
&\leq \mathbb{E}_{\mathbf{x}}\left[(\mathbf{x}^\top\mathbf{A}\mathbf{x})^2\right] \\
&= \langle \mathbf{A}, \mathcal{M}\mathbf{A}\rangle.
\end{aligned}$$

□

**Lemma 12.** *Under the hypothesis of Lemma 9, we have:*

$$\langle \mathbf{Z}, (\mathcal{I} - \gamma \mathcal{T}_b)\, \mathbf{Z} \rangle \geq 0 \; \forall \; \mathbf{Z} \in \mathcal{S}^+(\mathbb{R}^d).$$

*Proof of Lemma 12.* The first observation is that for any PSD matrix $\mathbf{Z} \in \mathcal{S}^+(\mathbb{R}^d)$, using lemma 11, $\langle \mathbf{Z}, \mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{Z} \rangle \leq \langle \mathbf{Z}, \mathcal{M}\mathbf{Z} \rangle$.

This implies that $\langle \mathbf{Z}, \mathcal{T}_b\mathbf{Z} \rangle = \langle \mathbf{Z}, (\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma}{b}\mathcal{M} - \gamma(\frac{b-1}{b})\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})\mathbf{Z} \rangle \leq \langle \mathbf{Z}, (\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \gamma\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})\mathbf{Z} \rangle$. The proof rounds up by employing Lemma 10. □

*Proof of Lemma 9.* The proof goes via a contradiction. Let us assume that there exists $\mathbf{Z} \in \mathcal{S}(d)$ such that the conclusion of Lemma 9 is false. So we have $\langle \mathbf{Z}, (\mathcal{I} - \gamma\mathcal{T}_b)\, \mathbf{Z} \rangle < -\langle \mathbf{Z}, \mathbf{Z} \rangle$, which implies that

$$\langle \mathbf{Z}, \mathcal{T}_b\mathbf{Z} \rangle > 2 \left\| \mathbf{Z} \right\|_F^2 / \gamma \tag{18}$$

Using the spectral decomposition theorem, let $\mathbf{Z} = \mathbf{Z}_1 - \mathbf{Z}_2$ where $\mathbf{Z}_1, \mathbf{Z}_2 \in \mathcal{S}^+(\mathbb{R}^d)$ and $\mathbf{Z}_1\mathbf{Z}_2 = \mathbf{Z}_2\mathbf{Z}_1 = 0$. Let $\widehat{Z} \stackrel{\text{def}}{=} \mathbf{Z}_1 + \mathbf{Z}_2$. This means that $\left\| \mathbf{Z} \right\|_F^2 = \left\| \widehat{Z} \right\|_F^2 = \left\| \mathbf{Z}_1 \right\|_F^2 + \left\| \mathbf{Z}_2 \right\|_F^2$. Since by hypothesis we have $\langle \widehat{Z}, \mathcal{T}_b\widehat{Z} \rangle > 0$, adding this to (18), we obtain

$$\begin{aligned}
&\langle Z, \mathcal{T}_b\mathbf{Z} \rangle + \langle \widehat{Z}, \mathcal{T}_b\widehat{Z} \rangle > 2 \left\| \mathbf{Z} \right\|_F^2 / \gamma \\
\Rightarrow\; & 2\langle \mathbf{Z}_1, \mathcal{T}_b\mathbf{Z}_1 \rangle + 2\langle \mathbf{Z}_2, \mathcal{T}_b\mathbf{Z}_2 \rangle > 2(\left\| \mathbf{Z}_1 \right\|_F^2 + \left\| \mathbf{Z}_2 \right\|_F^2)/\gamma \\
\Rightarrow\; & \text{either} \quad \langle \mathbf{Z}_1, \mathcal{T}_b\mathbf{Z}_1 \rangle > \left\| \mathbf{Z}_1 \right\|_F^2 / \gamma \quad \text{or} \quad \langle \mathbf{Z}_2, \mathcal{T}_b\mathbf{Z}_2 \rangle > \left\| \mathbf{Z}_2 \right\|_F^2 / \gamma \\
\Rightarrow\; & \text{either} \quad \langle \mathbf{Z}_1, (\mathcal{I} - \gamma\mathcal{T}_b)\mathbf{Z}_1 \rangle < 0 \quad \text{or} \quad \langle \mathbf{Z}_2, (\mathcal{I} - \gamma\mathcal{T}_b)\mathbf{Z}_2 \rangle < 0.
\end{aligned}$$

This contradicts with the conclusion of Lemma (12) which shows that for every $\mathbf{Z} \in \mathcal{S}^+(\mathbb{R}^d)$, we have $\langle \mathbf{Z}, (\mathcal{I} - \gamma\mathcal{T}_b)\mathbf{Z} \rangle \geq 0$. This proves that there does not exist $\mathbf{Z} \in \mathcal{S}(d)$ for which (18) holds, thus concluding the proof of Lemma 9.

Note that in dimension $d = 1$, we have:

$$\begin{aligned}
\mathcal{I} - \gamma\mathcal{T}_b &= 1 - \gamma(2\mathbb{E}\left[x^2\right] - \frac{\gamma}{b}\mathbb{E}\left[x^4\right] - \gamma(\frac{b-1}{b})\mathbb{E}\left[x^2\right]^2) \\
&= 1 - 2\gamma\mathbb{E}\left[x^2\right] + \frac{\gamma^2}{b}\mathbb{E}\left[x^4\right] + \gamma^2(\frac{b-1}{b})\mathbb{E}\left[x^2\right]^2 \\
&\geq 1 - 2\gamma\mathbb{E}\left[x^2\right] + \gamma^2\mathbb{E}\left[x^2\right]^2 \\
&= (1 - \gamma\mathbb{E}\left[x^2\right])^2 > 0
\end{aligned}$$

with the inequality relating the second to the fourth moment namely, $\mathbb{E}\left[x^2\right]^2 \leq \mathbb{E}\left[x^4\right]$ being through the application of Jensen's inequality.

□

**Lemma 13.** *With $\gamma > 0$, $\mathcal{T}_b = \mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma}{b}M - \gamma(\frac{b-1}{b})\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}$, we have $\forall \mathbf{A} \in \mathcal{S}(d)$, $2\,\mathrm{Tr}(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}\mathbb{E}(\mathbf{x}^T\mathbf{A}\mathbf{x})^2 - \gamma(\frac{b-1}{b})\,\mathrm{Tr}\;\mathbf{A}^T\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A} > 0$, then,*

- $\|\mathcal{I} - \gamma\mathcal{T}_b\| < 1$
- $\|\mathbf{I} - \gamma\mathbf{H}\| < 1$

*Proof.* This lemma generalizes the lemma governing contraction of all associated operators appearing in Défossez and Bach [1]. We note that the condition $\forall\, \mathbf{A} \in \mathcal{S}(d)$, $2\,\mathrm{Tr}(\mathbf{A}^T\mathbf{H}\mathbf{A}) - \frac{\gamma}{b}\mathbb{E}(\mathbf{x}^T\mathbf{A}\mathbf{x})^2 - \gamma(\frac{b-1}{b})\,\mathrm{Tr}\,\mathbf{A}^T\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}\mathbf{A} > 0$ implies $\mathcal{T}_b \succ 0$, with the step size set as $0 < \gamma < \gamma_{b,\max} < \frac{b}{b+1}(\frac{2}{\lambda_{\max}(\mathbf{H})}) < \frac{2}{\lambda_{\max}(\mathbf{H})}$.

Next, to prove $\|\mathcal{I} - \gamma\mathcal{T}_b\| < 1$, we require $-\mathcal{I} \prec \mathcal{I} - \gamma\mathcal{T}_b \prec \mathcal{I}$, the right hand side of which follows from the fact that $\mathcal{T}_b \succ 0$. The left hand side follows using lemma 9.

We know that $\mathbf{H} \succ 0$. Next, to ensure $\|\mathbf{I} - \gamma\mathbf{H}\| < 1$, we require $-\mathbf{I} \prec \mathbf{I} - \gamma\mathbf{H} \prec \mathbf{I}$. The right side of the inequality holds owing to the fact that $\mathbf{H} \succ 0$. The left side is equivalent to having $\mathbf{H} \prec \frac{2}{\gamma}\mathbf{I}$, or equivalently, $\gamma < \frac{2}{\lambda_{\max}(\mathbf{H})}$ which follows from corollary 8.1 derived from lemma 8.

□

### A.1.7 Proof of Lemma 7

We first note that with our choice of $\gamma_{b,\max}$, we have $\mathcal{T}_b \succeq 0$. This is because

$$\begin{aligned}
\langle \mathbf{A}, \mathcal{T}_b \mathbf{A} \rangle &= 2\,\mathrm{Tr}\left(\mathbf{AHA}\right) - \frac{\gamma}{b}\mathbb{E}\left[\langle \mathbf{x}, \mathbf{Ax}\rangle^2\right] - \frac{\gamma(b-1)}{b}\langle \mathbf{H}, \mathbf{AHA}\rangle \\
&\geq 2\,\mathrm{Tr}\left(\mathbf{AHA}\right) - \frac{\gamma}{b}\mathbb{E}\left[\|\mathbf{x}\|^2 \|\mathbf{Ax}\|^2\right] - \frac{\gamma(b-1)}{b}\|\mathbf{H}\|\,\mathrm{Tr}\left(\mathbf{AHA}\right) \\
&\geq 2\,\mathrm{Tr}\left(\mathbf{AHA}\right) - \frac{\gamma}{b}R^2\mathbb{E}\left[\|\mathbf{Ax}\|^2\right] - \frac{\gamma(b-1)}{b}\|\mathbf{H}\|\,\mathrm{Tr}\left(\mathbf{AHA}\right) \\
&\geq \left(2 - \frac{\gamma}{b}\left(R^2 + (b-1)\|\mathbf{H}\|\right)\right)\mathrm{Tr}\left(\mathbf{AHA}\right),
\end{aligned}$$

and using the definition of $\gamma_{b,\max}$ finishes the claim.

The rest of the proof is a generalization of the bound on the contraction factor provided by Défossez and Bach [1]. As we operate with a step size $0 < \gamma < \gamma_{b,\max}$, all the operators contract, i.e., $\rho_b < 1$ and $\mathcal{T}_b^{\gamma} \succ 0$ as per lemma 13. Furthermore, given the largest and smallest eigenvalues of $\mathbf{H}$ as $\lambda_{\max}(\mathbf{H})$ and $\lambda_{\min}(\mathbf{H})$, then, we denote the largest and smallest eigenvalues of $\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}}$ as $2\lambda_{\max}(\mathbf{H})$ and $2\lambda_{\min}(\mathbf{H})$.

We first provide bounds on the operator norm of $\mathcal{T}_b^{\gamma}$. Let $\lambda_{\max}(\mathcal{T}_b^{\gamma})$ and $\lambda_{\min}(\mathcal{T}_b^{\gamma})$ represent the largest and smallest eigenvalues of the operator $\mathcal{T}_b^{\gamma}$. Then,

$$\rho_{\mathcal{T}_b^{\gamma}} = \max\{1 - \gamma\lambda_{\min}(\mathcal{T}_b^{\gamma})), \gamma\lambda_{\max}(\mathcal{T}_b^{\gamma}) - 1\}$$

For a particular step size $\gamma$, we define $\alpha = \frac{\gamma}{\gamma_{b,\max}}$ implying the following:

$$\begin{aligned}
\mathcal{T}_b^{\gamma} &= \mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma}{b}\mathcal{M} - \gamma(\frac{b-1}{b})\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}} \\
&= (1-\alpha)(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}}) + \alpha(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}} - \frac{\gamma_{b,\max}}{b}\mathcal{M} - \gamma_{b,\max}(\frac{b-1}{b})\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}}) \\
&= (1-\alpha)(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}}) + \alpha\mathcal{T}_b^{\gamma_{b,\max}}
\end{aligned}$$

Since $\mathcal{T}_b^{\gamma_{b,\max}} \succ 0$, $\mathcal{T}_b^{\gamma} \succ 2(1-\alpha)\lambda_{\min}(\mathbf{H})\mathcal{I}$, implying $\lambda_{\min}(\mathcal{T}_b^{\gamma}) \geq 2(1-\alpha)\lambda_{\min}(\mathbf{H})$.

Next, $-\mathcal{I} \prec \mathcal{I} - \gamma_{b,\max}\mathcal{T}_b^{\gamma_{b,\max}}$ implying $\mathcal{T}_b^{\gamma_{b,\max}} \prec \frac{2}{\gamma_{b,\max}}\mathcal{I}$. This implies that,

$$\begin{aligned}
\mathcal{T}_b^{\gamma} &= (1-\alpha)(\mathcal{H}_{\mathcal{L}} + \mathcal{H}_{\mathcal{R}}) + \alpha\mathcal{T}_b^{\gamma_{b,\max}} \\
&\prec 2(1-\alpha)\lambda_{\max}(\mathbf{H})\mathcal{I} + \frac{2\alpha}{\gamma_{b,\max}}\mathcal{I}
\end{aligned}$$

implying $\lambda_{\max}(\mathcal{T}_b^{\gamma}) \leq 2(1-\alpha)\lambda_{\max}(\mathbf{H}) + \frac{2\alpha}{\gamma_{b,\max}}$. This when plugged into the operator norm of $\mathcal{T}_b^{\gamma}$ yields,

$$\begin{aligned}
\rho_{T_b^{\gamma}} &= \max\{1 - 2\alpha(1-\alpha)\gamma_{b,\max}\lambda_{\min}(\mathbf{H}), -1 + 2\alpha\gamma_{b,\max}((1-\alpha)\lambda_{\max}(\mathbf{H}) + \frac{\alpha}{\gamma_{b,\max}})\} \\
&= \max\{1 - 2\gamma(1 - \frac{\gamma}{\gamma_{b,\max}})\lambda_{\min}(\mathbf{H}), -1 + 2\gamma(1 - \frac{\gamma}{\gamma_{b,\max}})\lambda_{\max}(\mathbf{H}) + 2(\frac{\gamma}{\gamma_{b,\max}})^2\}
\end{aligned}$$

Similarly, the operator norm of $\mathbf{I} - \gamma\mathbf{H}$ is bounded as: $\rho_{Hb} = 1 - \gamma\lambda_{\min}(\mathbf{H})$, since in dimension $d \geq 2$, $1 - \gamma\lambda_{\min}(\mathbf{H}) \geq \gamma\lambda_{\max}(\mathbf{H}) - 1$. In d=1, $\rho_{Hb} = |1 - \gamma\lambda_{\min}(\mathbf{H})|$. This rounds up the claim in lemma 7.

### A.1.8 Proof of Theorem 1

The proof of Theorem 1 builds on the results of lemmas 7, 13 and Theorems 5, 6.

By instantiating Equation 4 of Theorem 5 and Equation 5 of Theorem 6 with $N = \frac{n}{b} - s$, and noting that the output of the Algorithm 1 is $\overline{\mathbf{w}}$, the result follows, i.e.:

$$\mathbb{E}\left[(\overline{\mathbf{w}} - \mathbf{w}^*)(\overline{\mathbf{w}} - \mathbf{w}^*)^{\top}\right] = \frac{1}{(\frac{n}{b} - s)^2\gamma^2}\mathcal{G}\mathcal{T}_b^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s\mathbf{E_0} + \mathcal{O}\left(\frac{\rho_b^{n/b}}{\gamma\lambda_{\min}(\mathbf{H})\left(\frac{n}{b} - s\right)}\|\mathbf{E_0}\|_F\right)$$

$$+\frac{1}{n-sb}\left[\mathcal{I}-\frac{1}{\gamma\left(\frac{n}{b}-s\right)}(\mathcal{H}_{\mathcal{L}}{}^{-1}+\mathcal{H}_{\mathcal{R}}{}^{-1})\right]\mathcal{G}\mathcal{T}_b{}^{-1}\mathbf{\Sigma}$$
$$+\frac{2}{\gamma b\left(\frac{n}{b}-s\right)^2}\mathcal{H}_{\mathcal{L}}{}^{-1}\mathcal{H}_{\mathcal{R}}{}^{-1}\mathcal{T}_b{}^{-1}\mathbf{\Sigma}+\mathcal{O}\left(\frac{\rho_b^{\frac{n}{b}-s+1}}{(n-bs)\lambda_{\min}(\mathbf{H})\lambda_{\min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F\right)$$

## A.2 Proofs of some useful Lemmas

**Lemma 14.**

$$\mathbb{E}\left[\|\mathbf{x}\|^2\mathbf{x}\mathbf{x}^\top\right]\preceq\frac{2}{\gamma_{1,max}}\cdot\mathbf{H}.$$

*Proof.* The proof of the claim above is straightforward given that $\mathbb{E}\|\mathbf{x}\|^2\mathbf{x}\otimes\mathbf{x}\preceq R^2\mathbf{H}$ and owing to the definition of $\gamma_{1,\max}=\frac{2}{R^2}$.

□

**Lemma 15.** *Denoting the assumption (A) $\gamma\leq\gamma_{b,max}/2$,*

1. *With (A) in place, $\mathcal{T}_b{}^{-1}\mathbf{W}\succeq 0$ for every $\mathbf{W}\in\mathcal{S}(d),\ \mathbf{W}\succeq 0$*
2. *With (A) in place, $\mathrm{Tr}\left(\mathcal{T}_b{}^{-1}\mathbf{W}\right)\leq\frac{2}{\lambda_{min}}\cdot\mathrm{Tr}\left(\mathbf{W}\right)\ \forall\ \mathbf{W}\in\mathcal{S}(d),\ \mathbf{W}\succeq 0$*
3. *With (A) in place, $\left\|\left(\mathcal{H}_{\mathcal{L}}{}^{-1}+\mathcal{H}_{\mathcal{R}}{}^{-1}-\gamma\mathcal{I}\right)\mathbf{W}\right\|_*\leq\frac{3}{\lambda_{min}}\cdot\|\mathbf{W}\|_*$ for every $\mathbf{W}\in\mathbb{R}^{d\times d}$*
4. *With (A) in place, $\mathrm{Tr}\left(\left(\mathcal{H}_{\mathcal{L}}{}^{-1}+\mathcal{H}_{\mathcal{R}}{}^{-1}-\gamma\mathcal{I}\right)\mathcal{T}_b{}^{-1}\mathbf{W}\right)\leq\frac{6}{\lambda_{min}^2}\cdot\mathrm{Tr}\left(\mathbf{W}\right)\ \forall\ \mathbf{W}\in\mathcal{S}(d),\ \mathbf{W}\succeq 0$*
5. *$\left(\mathcal{I}-\gamma\mathcal{T}_b\right)\mathbf{W}\succeq 0\ \forall\ \mathbf{W}\in\mathcal{S}(d),\ \mathbf{W}\succeq 0$*
6. *With (A) in place, $\mathrm{Tr}\left(\left(\mathcal{I}-\gamma\mathcal{T}_b\right)\mathbf{W}\right)\leq(1-\gamma\lambda_{min}/2)\,\mathrm{Tr}\left(\mathbf{W}\right)\ \forall\ \mathbf{W}\in\mathcal{S}(d),\ \mathbf{W}\succeq 0$*
7. *$\mathrm{Tr}\left(\left(\mathcal{H}_{\mathcal{R}}+\mathcal{H}_{\mathcal{L}}\right)^{-1}\mathbf{A}\right)=\frac{1}{2}\,\mathrm{Tr}\left(\mathbf{H}^{-1}\mathbf{A}\right)\ \forall\ \mathbf{A}\in\mathcal{S}^+(\mathbb{R}^d)$*
8. *With (A) in place,*

$$\mathrm{Tr}\left(\mathcal{T}_b{}^{-1}\mathbf{\Sigma}\right)\leq\frac{3}{2}\,\mathrm{Tr}\left(\mathbf{H}^{-1}\mathbf{\Sigma}\right)$$

*Proof. Proof of claim 1 in Lemma 15*: We require to prove $\mathcal{T}_b{}^{-1}$ operating on a PSD matrix produces a PSD matrix, or in other words, $\mathcal{T}_b{}^{-1}$ is a PSD map.

$$\begin{aligned}\mathcal{T}_b{}^{-1}&=[\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}}-\frac{\gamma}{b}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})]^{-1}\\&=(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{\frac{1}{2}}[\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}}-\frac{\gamma}{b}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})]^{-1}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{\frac{1}{2}}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}\\&=(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}[\mathcal{I}-\frac{\gamma}{b}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}]^{-1}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}\end{aligned}\tag{19}$$

Now, we prove that $\|\frac{\gamma}{b}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}\|<1$. Given $\gamma<\gamma_{b,\max}/2$, we employ lemma 8 to note that $\mathcal{T}_b\succ 0$.

$$\begin{aligned}&\mathcal{T}_b\succ 0\\&\Rightarrow\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}}-\frac{\gamma}{b}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})\succ 0\\&\Rightarrow\frac{\gamma}{b}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})\prec\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}}\\&\Rightarrow\frac{\gamma}{b}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}\prec\mathcal{I}\\&\Rightarrow\|\frac{\gamma}{b}(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}(\mathcal{M}+(b-1)\mathcal{H}_{\mathcal{L}}\mathcal{H}_{\mathcal{R}})(\mathcal{H}_{\mathcal{L}}+\mathcal{H}_{\mathcal{R}})^{-\frac{1}{2}}\|<1\end{aligned}$$

With this fact in place, we employ Taylor series to expand ${\mathcal{T}_b}^{-1}$ in equation 19, i.e.:

$$\begin{aligned}{\mathcal{T}_b}^{-1} &= (\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-\frac{1}{2}}\sum_{i=0}^{\infty}(\frac{\gamma}{b}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-\frac{1}{2}}(\mathcal{M}+(b-1)\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R})(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-\frac{1}{2}})^i(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-\frac{1}{2}} \\ &= \sum_{i=0}^{\infty}(\frac{\gamma}{b}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}(\mathcal{M}+(b-1)\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R}))^i(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\end{aligned}$$

The proof completes by employing the following facts: Using Lyapunov's theorem, we know $(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}$ is a PSD map, i.e. if $(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}(A)=B$, then, if $A$ is PSD $\implies$ $B$ is PSD. Furthermore, $\mathcal{M}$ is also a PSD map, i.e. if $A_1$ is PSD, $\mathcal{M}(A_1)=E[(x^TA_1x)x\otimes x]$ is PSD as well. Finally, $\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R}$ is also a PSD map, since, if $A_2$ is PSD, then, $\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R}(A_2)=HA_2H$ which is PSD as well. With all these facts in place, we note that each term in the Taylor's expansion above is a PSD map implying the overall map is PSD as well, thus rounding up the proof to claim 1 in Lemma 15.

*Proof of claim 2 in Lemma 15*: Since we know that ${\mathcal{T}_b}^{-1}$ is a PSD map, it suffices to prove the following stronger statement where we replace the condition $\mathbf{W}\succeq 0$ with ${\mathcal{T}_b}^{-1}\mathbf{W}\succeq 0$.

$$\operatorname{Tr}\left({\mathcal{T}_b}^{-1}\mathbf{W}\right)\leq\frac{1}{\lambda_{\min}}\cdot\operatorname{Tr}(\mathbf{W})\ \ \forall\ \mathbf{W}\in\mathcal{S}(d),\ {\mathcal{T}_b}^{-1}\mathbf{W}\succeq 0.$$

This is equivalent to showing that

$$\operatorname{Tr}(\mathbf{U})\leq\frac{1}{\lambda_{\min}}\cdot\operatorname{Tr}(\mathcal{T}_b\mathbf{U})\ \ \forall\ \mathbf{U}\in\mathcal{S}(d),\ \mathbf{U}\succeq 0. \tag{20}$$

Since $\mathcal{T}_b\mathbf{U}=\mathbf{H}\mathbf{U}+\mathbf{U}\mathbf{H}-\frac{\gamma}{b}\mathcal{M}\mathbf{U}-\frac{\gamma(b-1)}{b}\mathbf{H}\mathbf{U}\mathbf{H}$, we bound the following quantities:

$$\operatorname{Tr}(\mathcal{M}\mathbf{U})=\mathbb{E}\left[\|\mathbf{x}\|^2\mathbf{x}^\top\mathbf{U}\mathbf{x}\right]=\langle\mathbb{E}\left[\|\mathbf{x}\|^2\mathbf{x}\mathbf{x}^\top\right],\mathbf{U}\rangle\leq\frac{2}{\gamma_{1,\max}}\langle\mathbf{H},\mathbf{U}\rangle, \tag{21}$$

where the last step follows from Lemma 14 and the fact that $\mathbf{U}\succeq 0$. We also have

$$\operatorname{Tr}(\mathbf{H}\mathbf{U}\mathbf{H})=\langle\mathbf{H},\mathbf{H}^{1/2}\mathbf{U}\mathbf{H}^{1/2}\rangle\leq\|\mathbf{H}\|_2\operatorname{Tr}\left(\mathbf{H}^{1/2}\mathbf{U}\mathbf{H}^{1/2}\right)=\|\mathbf{H}\|_2\operatorname{Tr}(\mathbf{H}\mathbf{U}). \tag{22}$$

Using (21) and (22), we have

$$\begin{aligned}\operatorname{Tr}(\mathcal{T}_b\mathbf{U}) &= 2\operatorname{Tr}(\mathbf{H}\mathbf{U})-\frac{\gamma}{b}\operatorname{Tr}(\mathcal{M}\mathbf{U})-\frac{\gamma(b-1)}{b}\operatorname{Tr}(\mathbf{H}\mathbf{U}\mathbf{H}) \\ &\geq 2\operatorname{Tr}(\mathbf{H}\mathbf{U})-\frac{\gamma}{b\gamma_{1,\max}}\operatorname{Tr}(\mathbf{H}\mathbf{U})-\frac{\gamma\|\mathbf{H}\|_2(b-1)}{b}\operatorname{Tr}(\mathbf{H}\mathbf{U}) \\ &= \left(2-\frac{\gamma}{b\gamma_{1,\max}}-\frac{\gamma\|\mathbf{H}\|_2(b-1)}{b}\right)\operatorname{Tr}(\mathbf{H}\mathbf{U}).\end{aligned}$$

Since $\frac{\gamma}{b\gamma_{1,\max}}<\frac{\gamma_{b,\max}}{2b\gamma_{1,\max}}\leq\frac{1}{2}$, and $\gamma\|\mathbf{H}\|_2\leq\gamma_{b,\max}\cdot\|\mathbf{H}\|_2/2\leq\frac{2}{\|\mathbf{H}\|_2}\cdot\|\mathbf{H}\|_2/2=1$, we have:

$$\operatorname{Tr}(\mathcal{T}_b\mathbf{U})\geq\frac{1}{2}\operatorname{Tr}(\mathbf{H}\mathbf{U})\geq\frac{\lambda_{\min}}{2}\cdot\operatorname{Tr}(\mathbf{U}).$$

This proves the claim.

*Proof of claim 3 in Lemma 15*: Again using Lemma 16, we know that

$$\begin{aligned}&\left\|\left({\mathcal{H}_\mathcal{L}}^{-1}+{\mathcal{H}_\mathcal{R}}^{-1}-\gamma\mathcal{I}\right)\mathbf{W}\right\|_* = \left\|\mathbf{H}^{-1}\mathbf{W}+\mathbf{W}\mathbf{H}^{-1}-\gamma\mathbf{W}\right\|_* \\ &\leq\left\|\mathbf{H}^{-1}\mathbf{W}\right\|_*+\left\|\mathbf{W}\mathbf{H}^{-1}\right\|_*+\gamma\|\mathbf{W}\|_*\leq\left(\frac{2}{\lambda_{\min}}+\gamma\right)\|\mathbf{W}\|_*\leq\frac{3}{\lambda_{\min}}\|\mathbf{W}\|_*,\end{aligned}$$

where we used the fact that $\gamma<\gamma_{b,\max}/2\leq\frac{1}{\|\mathbf{H}\|_2}$.

*Proof of claim 4 in Lemma 15*: This follows easily from combining Claims 1, 2 and 3.

$$\operatorname{Tr}\left(\left({\mathcal{H}_\mathcal{L}}^{-1}+{\mathcal{H}_\mathcal{R}}^{-1}-\gamma\mathcal{I}\right){\mathcal{T}_b}^{-1}\mathbf{W}\right)\leq\left\|\left({\mathcal{H}_\mathcal{L}}^{-1}+{\mathcal{H}_\mathcal{R}}^{-1}-\gamma\mathcal{I}\right){\mathcal{T}_b}^{-1}\mathbf{W}\right\|_*\ (\text{ from Lemma 16})$$

$$
\begin{aligned}
&\leq \frac{3}{\lambda_{\min}} \cdot \left\|\mathcal{T}_b^{-1}\mathbf{W}\right\|_* \text{ (from Claim 3)} \\
&= \frac{3}{\lambda_{\min}} \cdot \operatorname{Tr}\left(\mathcal{T}_b^{-1}\mathbf{W}\right) \text{ (from Claim 1)} \\
&\leq \frac{6}{\lambda_{\min}^2} \cdot \operatorname{Tr}\left(\mathbf{W}\right) \text{ (from Claim 2).}
\end{aligned}
$$

*Proof of claim 5 in Lemma 15*: Since $(\mathcal{I} - \gamma\mathcal{T}_b)\,\mathbf{W} = (\mathbf{I} - \gamma\mathbf{H})\,\mathbf{W}\,(\mathbf{I} - \gamma\mathbf{H}) + \gamma^2/b\,(\mathcal{M}\mathbf{W} - \mathbf{HWH})$, it suffices to show that $(\mathcal{M}\mathbf{W} - \mathbf{HWH}) \succeq 0$. So we need to show

$$
\begin{aligned}
&(\mathcal{M}\mathbf{W} - \mathbf{HWH}) \succeq 0 \\
\Leftrightarrow &\mathcal{M}\mathbf{W} \succeq \mathbf{HWH} \\
\Leftrightarrow &\mathbf{u}^\top(\mathcal{M}\mathbf{W})\mathbf{u} \geq \mathbf{u}^\top\mathbf{HWHu}\ \forall\ \mathbf{u} \in \mathbb{R}^d \\
\Leftrightarrow &\mathbb{E}_x\left[(\mathbf{u}^\top\mathbf{x})^2(\mathbf{x}^\top\mathbf{Wx})\right] \geq \mathbb{E}_{\mathbf{x}_1,\mathbf{x}_2}\left[(\mathbf{x}_1{}^\top\mathbf{W}\mathbf{x}_2)(\mathbf{u}^\top\mathbf{x}_1)(\mathbf{u}^\top\mathbf{x}_2)\right]\ \forall\ \mathbf{u} \in \mathbb{R}^d.
\end{aligned}
$$

This last statement is proved using Cauchy-Schwarz inequality as follows:

$$
\begin{aligned}
&\mathbb{E}\left[(\mathbf{x}_1{}^\top\mathbf{W}\mathbf{x}_2)(\mathbf{u}^\top\mathbf{x}_1)(\mathbf{u}^\top\mathbf{x}_2)\right] \\
&\leq \mathbb{E}\left[\left\|\mathbf{W}^{1/2}\mathbf{x}_1\right\|\left\|\mathbf{W}^{1/2}\mathbf{x}_2\right\|\left|\mathbf{u}^\top\mathbf{x}_1\right|\left|\mathbf{u}^\top\mathbf{x}_2\right|\right] \\
&\leq \left(\mathbb{E}\left[(\mathbf{x}_1{}^\top\mathbf{W}\mathbf{x}_1)(\mathbf{x}_2{}^\top\mathbf{W}\mathbf{x}_2)(\mathbf{u}^\top\mathbf{x}_1)^2(\mathbf{u}^\top\mathbf{x}_2)^2\right]\right)^{1/2} = \mathbb{E}\left[(\mathbf{u}^\top\mathbf{x})^2(\mathbf{x}^\top\mathbf{Wx})\right].
\end{aligned}
$$

This finishes the proof.

*Proof of claim 6 in Lemma 15*: Since in the proof of Claim 2 we proved (20) which says that

$$
\operatorname{Tr}\left(\mathcal{T}_b\mathbf{U}\right) \geq \lambda_{\min} \cdot \operatorname{Tr}\left(\mathbf{U}\right)\ \forall\ \mathbf{U} \in \mathcal{S}(d),\ \mathbf{U} \succeq 0,
$$

Claim 6 follows.

*Proof of claim 7 in Lemma 15*:

We know that the operator $\left(\mathcal{H}_\mathcal{R} + \mathcal{H}_\mathcal{L}\right)^{-1}$ is a PSD map, i.e, it maps PSD matrices to PSD matrices. Since $\mathbf{A} \succeq 0$, we replace this condition with $\mathbf{U} = \left(\mathcal{H}_\mathcal{R} + \mathcal{H}_\mathcal{L}\right)^{-1}\mathbf{A} \succeq\ 0$ implying, we need to show the following:

$$
\operatorname{Tr}\left(\mathbf{U}\right) = \frac{1}{2}\operatorname{Tr}\left(\mathbf{H}^{-1}\mathbf{U}\right)\ \forall\ \mathbf{U} \succeq 0
$$

Examining the right hand side, we see the following:

$$
\begin{aligned}
\frac{1}{2}\operatorname{Tr}\left(\mathbf{H}^{-1}\mathbf{U}\right) &= \frac{1}{2}\operatorname{Tr}\left(\mathbf{H}^{-1}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})\mathbf{U}\right) \\
&= \frac{1}{2}\operatorname{Tr}\left(\mathbf{H}^{-1}\mathbf{HU} + \mathbf{H}^{-1}\mathbf{UH}\right) \\
&= \operatorname{Tr}\left(\mathbf{U}\right)
\end{aligned}
$$

thus wrapping up the proof of claim 8.

*Proof of claim 8 in Lemma 15*:

We first denote $\mathbf{U} = \mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R} - \gamma\frac{b-1}{b}\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R}$. Then, we note that the taylor's expansion holds since $\gamma < \gamma_{b,\max}$ and hence $\mathcal{T}_b \succ 0$.

$$
\begin{aligned}
\mathcal{T}_b{}^{-1}\mathbf{\Sigma} &= \left(\mathbf{U} - \frac{\gamma}{b}\mathcal{M}\right)^{-1}\mathbf{\Sigma} \\
&= \mathbf{U}^{-1/2}(\mathcal{I} - \frac{\gamma}{b}\mathbf{U}^{-1/2}\mathcal{M}\mathbf{U}^{-1/2})^{-1}\mathbf{U}^{-1/2}\mathbf{\Sigma} \\
&= \mathbf{U}^{-1}\mathbf{\Sigma} + \sum_{i=1}^{\infty}(\frac{\gamma}{b})^i(\mathbf{U}^{-1}\mathcal{M})^i(\mathbf{U}^{-1}\mathbf{\Sigma})
\end{aligned} \tag{23}
$$

Before we proceed further, we provide some upper bounds on the spectral properties of the map $\mathbf{U}^{-1}$ as applied to any psd matrix $\mathbf{B} \succeq 0$:

$$
\mathbf{U}^{-1}\mathbf{B} = \left(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R} - \gamma\frac{b-1}{b}\mathcal{H}_\mathcal{L}\mathcal{H}_\mathcal{R}\right)^{-1}\mathbf{B}
$$

$$= (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1/2} \left( \mathcal{I} - \gamma \frac{b-1}{b} (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1/2} \mathcal{H}_\mathcal{L} \mathcal{H}_\mathcal{R} (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1/2} \right)^{-1} (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1/2} \mathbf{B}$$
$$= \sum_{i=0}^{\infty} \left( \gamma \frac{b-1}{b} (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathcal{H}_\mathcal{L} \mathcal{H}_\mathcal{R} \right)^i \hat{\mathbf{B}}$$

where, $\hat{\mathbf{B}} = (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B}$. Next,

$$\begin{aligned}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathcal{H}_\mathcal{L} \mathcal{H}_\mathcal{R} \hat{\mathbf{B}} &= (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{H} \hat{\mathbf{B}} \mathbf{H} \\ &= \frac{1}{2} \hat{\mathbf{B}} \mathbf{H} \\ &\preceq \frac{\|\mathbf{H}\|_2}{2} \hat{\mathbf{B}}\end{aligned}$$

implying,

$$\begin{aligned}\mathbf{U}^{-1} \mathbf{B} &\preceq \sum_{i=0}^{\infty} \left( (\frac{\gamma}{2})(\frac{b-1}{b}) \|\mathbf{H}\|_2 \right)^i \hat{\mathbf{B}} \\ &= \left( \frac{1}{1 - \gamma \frac{b-1}{2b} \|\mathbf{H}\|_2} \right) \hat{\mathbf{B}} \quad , \gamma < \frac{2b}{b-1} \|\mathbf{H}\|_2\end{aligned}$$

Next, we establish some lower bounds on the spectral properties of the map $\mathbf{U}^{-1}$ when applied to any psd matrix $\mathbf{B} \succeq 0$. In particular, our main claim is:

$$(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B} \preceq \mathbf{U}^{-1} \mathbf{B}$$

In order to prove this claim, let us consider $\mathbf{C}_1 = (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B}$, and $\mathbf{C}_2 = \mathbf{U}^{-1} \mathbf{B}$. Given that $B \succeq 0$, we know that $\mathbf{C}_1 \succeq 0$ (by Lyapunov's Theorem), and $\mathbf{C}_2 \succeq 0$ (using an argument similar to claim 1, with the step size $\gamma < \frac{2b}{(b-1)\|\mathbf{H}\|_2}$ chosen in order to have $\mathbf{U} \succ 0$ for the arguments similar to claim 1 to hold). Then, we establish the claim using the following series of equations:

$$\begin{aligned}\mathbf{B} = (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R}) \mathbf{C}_1 &= (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R} - \gamma \frac{b-1}{b} \mathcal{H}_\mathcal{L} \mathcal{H}_\mathcal{R}) \mathbf{C}_2 \\ \mathbf{H}\mathbf{C}_1 + \mathbf{C}_1 \mathbf{H} &= \mathbf{H}\mathbf{C}_2 + \mathbf{C}_2 \mathbf{H} - \gamma \frac{b-1}{b} \mathbf{H}\mathbf{C}_2\mathbf{H} \\ \gamma \frac{b-1}{b} \mathbf{H}\mathbf{C}_2\mathbf{H} &= \mathbf{H}(\mathbf{C}_2 - \mathbf{C}_1) + (\mathbf{C}_2 - \mathbf{C}_1)\mathbf{H} \\ \gamma \frac{b-1}{b} \mathbf{H}\mathbf{C}_2\mathbf{H} &= (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})(\mathbf{C}_2 - \mathbf{C}_1) \\ \gamma \frac{b-1}{b} (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{H}\mathbf{C}_2\mathbf{H} &= \mathbf{C}_2 - \mathbf{C}_1\end{aligned}$$

We know that the left hand side is psd (using Lyapunov's theorem), and hence, $\mathbf{C}_2 \succeq \mathbf{C}_1$ which establishes the claim. Now, we combine the resulting lower bounds with the upper bounds to obtain:

$$\operatorname{Tr}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B} \leq \operatorname{Tr} \mathbf{U}^{-1} \mathbf{B} \leq \{ \frac{1}{1 - \frac{\gamma(b-1)\|\mathbf{H}\|_2}{2b}} \} \operatorname{Tr}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B} \tag{24}$$

$$\left\| (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B} \right\|_2 \leq \left\| \mathbf{U}^{-1} \mathbf{B} \right\|_2 \leq \{ \frac{1}{1 - \frac{\gamma(b-1)\|\mathbf{H}\|_2}{2b}} \} \left\| (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1} \mathbf{B} \right\|_2 \tag{25}$$

with the condition that $\gamma < \frac{2b}{(b-1)\|\mathbf{H}\|_2}$. Next, denoting $\mathbf{A} = \mathbf{U}^{-1} \mathbf{\Sigma}$, we have by equation 23:

$$\mathcal{T}_b{}^{-1} \mathbf{\Sigma} = \mathbf{A} + \sum_{i=1}^{\infty} (\mathbf{U}^{-1} \mathcal{M})^i \mathbf{A} \tag{26}$$

With this in place, we resort to bounding the behavior of equation 26 in a psd sense, with the following arguments:

$$\begin{aligned}\mathbf{U}^{-1}\mathcal{M}\mathbf{A} &= \mathbf{U}^{-1}\mathbb{E}(x^t\mathbf{A}x)x\otimes x\\ &\preceq \|\mathbf{A}\|_2 R^2\mathbf{U}^{-1}\mathbf{H}\\ &= \|\mathbf{A}\|_2 R^2\mathbf{A}'\end{aligned}$$

where, $\mathbf{A}' = \mathbf{U}^{-1}\mathbf{H}$. We note that $\mathbf{U}^{-1}$ is a psd map (following arguments similar to claim 1 in lemma 15 with the step size $\gamma < \frac{2b}{(b-1)\|\mathbf{H}\|_2}$), which is necessary for the above sequence of inequalities to hold. This implies,

$$\begin{aligned}(\mathbf{U}^{-1}\mathcal{M})^2\mathbf{A} &= \mathbf{U}^{-1}\mathcal{M}(\mathbf{U}^{-1}\mathcal{M}\mathbf{A})\\ &\preceq \|\mathbf{A}\|_2 R^2\mathbf{U}^{-1}\mathcal{M}\mathbf{A}'\\ &\preceq \|\mathbf{A}\|_2 R^2(\|\mathbf{A}'\|_2 R^2)\mathbf{A}'\end{aligned}$$

implying, the $i^{th}$ term in equation 26 is bounded as:

$$(U^{-1}\mathcal{M})^i\mathbf{A} \preceq \|\mathbf{A}\|_2 R^2(\|\mathbf{A}'\|_2 R^2)^{i-1}\mathbf{A}'$$

Using these bounds in equation 26 and summing up the resulting geometric series and applying trace on both sides, we get the following bound:

$$\operatorname{Tr}\mathcal{T}_b^{-1}\mathbf{\Sigma} \le \operatorname{Tr}\mathbf{U}^{-1}\mathbf{\Sigma} + \frac{\frac{\gamma}{b}\left\|\mathbf{U}^{-1}\mathbf{\Sigma}\right\|_2 R^2\operatorname{Tr}\mathbf{U}^{-1}\mathbf{H}}{1-\left(\frac{\gamma}{b}\|\mathbf{U}^{-1}\mathbf{H}\|_2 R^2\right)} \quad ,\gamma < \frac{b}{\|\mathbf{U}^{-1}\mathbf{H}\|_2 R^2} \tag{27}$$

To bound the behavior of the first term in equation 27, we employ the bound in equation 24 (which ensures statistically optimal rates), and is satisfied if $\gamma < \frac{2b}{(b-1)\|\mathbf{H}\|_2}$. Next, the denominator of the second term in equation 27 holds iff the geometric series converges, which is if $\gamma < \frac{b}{R^2\|\mathbf{U}^{-1}\mathbf{H}\|_2}$, which by employing equation 25 is upperbounded as $\gamma < \frac{b}{R^2\|(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{H}\|_2} = \frac{2b}{R^2}$. Finally, for the numerator of the second term in equation 27 to be statistically optimal, we require $\gamma < \frac{b\operatorname{Tr}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}}{\|\mathbf{U}^{-1}\mathbf{\Sigma}\|_2 R^2\operatorname{Tr}\mathbf{U}^{-1}\mathbf{H}}$, which by employing the bounds in equations 25 and 24, we obtain a bound as $\gamma < \frac{2b\operatorname{Tr}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}}{\|(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\|_2 R^2 d}$. Combining these three bounds, the step size that suffices for the statistically optimal rates are:

$$\gamma < \min\left(\frac{2b}{(b-1)\|\mathbf{H}\|_2}, \frac{2b}{R^2}, \frac{2b\operatorname{Tr}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}}{\|(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\|_2 R^2 d}\right)$$

We also note that the third term is always $\le$ the second term, implying,

$$\gamma < \min\left(\frac{2b}{(b-1)\|\mathbf{H}\|_2}, \frac{2b\operatorname{Tr}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}}{\|(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\|_2 R^2 d}\right)$$

In particular, with the choice of step size that is half of this maximum, by employing claim 7, we have

$$\operatorname{Tr}\mathcal{T}_b^{-1}\mathbf{\Sigma} \le \frac{3}{2}\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}$$

□

**Lemma 16.** *Recall that the nuclear norm $\|\mathbf{U}\|_*$ of a matrix $\mathbf{U}$ is defined to be the sum of its singular values. The following properties are satisfied by the nuclear norm:*

- $\|\mathbf{U}\mathbf{A}\|_* \le \|\mathbf{A}\|_2\|\mathbf{U}\|_*$ *for all* $\mathbf{U},\mathbf{A}\in\mathbb{R}^{d\times d}$.
- $\|\mathbf{U}\|_* = \operatorname{Tr}(\mathbf{U})$ *for all PSD matrices* $\mathbf{U}$.
- $\|\mathbf{U}\|_* = \max_{\substack{\mathbf{A}\in\mathbb{R}^{d\times d}\\ \|\mathbf{A}\|_2\le 1}}\langle\mathbf{U},\mathbf{A}\rangle$ *for all* $\mathbf{U}\in\mathbb{R}^{d\times d}$. *Moreover, the optimizer* $\mathbf{A}^*$ *satisfies* $\mathbf{A}^*\mathbf{U} = \mathbf{U}\mathbf{A}^* \succeq 0$.

## A.3 Proofs of Theorems 2 and 3

### A.3.1 Proof of Theorem 2

*Proof.* Since $\gamma \leq \frac{\gamma_{b,\max}}{3}$, we first invoke lemma 7 to get the contraction factor $\rho_b$:

$$\rho_b \leq \max(1 - 2\gamma(1 - \frac{\gamma}{\gamma_{b,\max}})\lambda_{\min}(\mathbf{H}), -1 + 2\gamma(1 - \frac{\gamma}{\gamma_{b,\max}})\lambda_{\max}(\mathbf{H}) + 2(\frac{\gamma}{\gamma_{b,\max}})^2, 1 - \gamma\lambda_{\min}(\mathbf{H})) \tag{28}$$

We now wish to get a finer bound on the contraction factor, for which, we first define $\alpha = \frac{\gamma}{\gamma_{b,\max}}$ in the equation 28 we subtract the second term from the first, to yield:

$$\begin{aligned}
&1 - 2\gamma(1-\alpha)\lambda_{\min}(\mathbf{H}) - (-1 + 2\gamma(1-\alpha)\lambda_{\max}(\mathbf{H}) + 2\alpha^2) \\
&= 2 - 2\alpha\gamma_{b,\max}(1-\alpha)(\lambda_{\min}(\mathbf{H}) + \lambda_{\max}(\mathbf{H})) - 2\alpha^2 \\
&\geq 2 - 4\alpha\gamma_{b,\max}(1-\alpha)\lambda_{\max}(\mathbf{H}) - 2\alpha^2 \\
&= 2(1 - 2\alpha(1-\alpha)\gamma_{b,\max}\lambda_{\max}(\mathbf{H}) - \alpha^2) \\
&\geq 2(1 - 4\alpha(1-\alpha) - \alpha^2) \\
&= 2(1 - 4\alpha + 3\alpha^2) \\
&= 2(1-\alpha)(1-3\alpha) \\
&\geq 0
\end{aligned}$$

implies that $\alpha \leq 1$ and $\alpha \leq \frac{1}{3}$, which is indeed true from the statement of the theorem. Once this is indeed true, we bound the contraction factor from equation 28 as:

$$\rho_b \leq \max(1 - 2\gamma(1-\alpha)\lambda_{\min}(\mathbf{H}), 1 - \gamma\lambda_{\min}(\mathbf{H}))$$

By setting $\alpha = \frac{1}{6}$, we get,

$$\rho_b \leq \max(1 - \frac{5}{9\kappa}, 1 - \frac{1}{3\kappa}) = 1 - \frac{1}{3\kappa}$$

With a bound on the contraction factor established explicitly, we now turn towards providing a scalar bound by taking the innerproduct of each term in equation 2 with $\mathbf{H}$, which helps in establishing the statement of theorem 2.

$$\begin{aligned}
\langle \mathbf{H}, \mathfrak{T}_1 \rangle &= \frac{1}{(\frac{n}{b} - s)^2\gamma^2} \langle \mathbf{H}, \mathcal{G}\mathcal{T}_b^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s \mathbf{E_0} \rangle \\
&= \frac{1}{(\frac{n}{b} - s)^2\gamma^2} \langle \mathcal{T}_b^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s \mathbf{E_0}, (\mathcal{H}_{\mathcal{L}}^{-1} + \mathcal{H}_{\mathcal{R}}^{-1} - \gamma\mathcal{I})\mathbf{H} \rangle \\
&\leq \frac{2}{(\frac{n}{b} - s)^2\gamma^2} \langle \mathcal{T}_b^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s \mathbf{E_0}, \mathbf{I} \rangle \\
&= \frac{2}{(\frac{n}{b} - s)^2\gamma^2} \operatorname{Tr}\left(\mathcal{T}_b^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s \mathbf{E_0}\right) \\
&\leq \frac{2\sqrt{d}}{(\frac{n}{b} - s)^2\gamma^2} \left\| \mathcal{T}_b^{-1}(\mathcal{I} - \gamma\mathcal{T}_b)^s \mathbf{E_0} \right\|_F \\
&\leq \frac{2\sqrt{d}}{(\frac{n}{b} - s)^2\gamma^2} \left\| \mathcal{T}_b^{-1} \right\| \left\| \mathcal{I} - \gamma\mathcal{T}_b \right\|^s \left\| \mathbf{E_0} \right\|_F \\
&\leq \frac{2\sqrt{d}}{(\frac{n}{b} - s)^2\gamma^2} \mathrm{e}^{(-\frac{s}{3\kappa})} \|\mathcal{T}_b^{-1}\| \|\mathbf{E_0}\|_F \\
&\leq \frac{2\sqrt{d}}{(\frac{n}{b} - s)^2\gamma^2} \mathrm{e}^{(-\frac{s}{3\kappa})} \|\mathcal{T}_b^{-1}\| \operatorname{Tr}(\mathbf{E_0}) \\
&\leq \frac{2\sqrt{d}}{\gamma^2(\frac{n}{b} - s)^2} \mathrm{e}^{(-\frac{s}{3\kappa})} \|\mathcal{T}_b^{-1}\| \left( \frac{1}{\lambda_{\min}(\mathbf{H})} \langle \mathbf{H}, \mathbf{E_0} \rangle \right)
\end{aligned}$$

$$\leq \frac{4\sqrt{d}}{\gamma^2(\frac{n}{b}-s)^2}\mathrm{e}^{(-\frac{s}{3\kappa})}\|\mathbf{H}^{-1}\|\left(\frac{1}{\lambda_{\min}(\mathbf{H})}\langle\mathbf{H},\mathbf{E_0}\rangle\right)$$
$$= \frac{4\sqrt{d}}{\gamma^2(\frac{n}{b}-s)^2(\lambda_{\min}(\mathbf{H}))^2}\mathrm{e}^{(-\frac{s}{3\kappa})}\langle\mathbf{H},\mathbf{E_0}\rangle$$

which wraps up bounding $\langle\mathbf{H},\mathfrak{T}_1\rangle$ by plugging in the value of $\gamma$. We switch our attention to the second term.

$$\begin{aligned}\langle\mathbf{H},\mathfrak{T}_2\rangle &= \|\mathbf{H}\|\,\mathrm{Tr}\left(\mathfrak{T}_2\right)\\ &\leq \sqrt{d}\,\|\mathbf{H}\|\,\|\mathfrak{T}_2\|_F\\ &\leq \sqrt{d}\,\|\mathbf{H}\|\cdot\frac{\rho_b^{n/b}}{\gamma\lambda_{\min}(\mathbf{H})\left(\frac{n}{b}-s\right)}\|\mathbf{E_0}\|_F\\ &\leq \frac{3\exp(-n/3b\kappa)\kappa^2\sqrt{d}}{(\frac{n}{b}-s)}\langle\mathbf{H},\mathbf{E}_0\rangle,\end{aligned}$$

which bounds the second term. Next, we proceed to bounding the third term:

$$\begin{aligned}\langle\mathbf{H},\mathfrak{T}_3\rangle &= \frac{1}{n-sb}\langle\mathbf{H},\left[\mathcal{I}-\frac{1}{\gamma\left(\frac{n}{b}-s\right)}({\mathcal{H}_{\mathcal{L}}}^{-1}+{\mathcal{H}_{\mathcal{R}}}^{-1})\right]\mathcal{G}{\mathcal{T}_b}^{-1}\mathbf{\Sigma}\rangle\\ &= \frac{1}{n-sb}\langle\mathbf{\Sigma},{\mathcal{T}_b}^{-1}\mathcal{G}\left[\mathcal{I}-\frac{1}{\gamma\left(\frac{n}{b}-s\right)}({\mathcal{H}_{\mathcal{L}}}^{-1}+{\mathcal{H}_{\mathcal{R}}}^{-1})\right]\mathbf{H}\rangle\\ &\leq \frac{1}{n-sb}\langle\mathbf{\Sigma},{\mathcal{T}_b}^{-1}({\mathcal{H}_{\mathcal{L}}}^{-1}+{\mathcal{H}_{\mathcal{R}}}^{-1}-\gamma\mathcal{I})\mathbf{H}\rangle\\ &\leq \frac{2}{n-sb}\langle\mathbf{\Sigma},{\mathcal{T}_b}^{-1}\mathbf{I}\rangle\\ &= \frac{2}{n-sb}\langle\mathbf{I},{\mathcal{T}_b}^{-1}\mathbf{\Sigma}\rangle\\ &\leq \frac{3\,\mathrm{Tr}\,\mathbf{H}^{-1}\mathbf{\Sigma}}{n-sb}\end{aligned}\tag{29}$$

Next, we proceed to the fourth term:

$$\begin{aligned}\langle\mathbf{H},\mathfrak{T}_4\rangle &= \frac{2}{\gamma b\left(\frac{n}{b}-s\right)^2}\langle\mathbf{H},{\mathcal{H}_{\mathcal{L}}}^{-1}{\mathcal{H}_{\mathcal{R}}}^{-1}{\mathcal{T}_b}^{-1}\mathbf{\Sigma}\rangle\\ &= \frac{2}{\gamma b\left(\frac{n}{b}-s\right)^2}\langle\mathbf{\Sigma},{\mathcal{T}_b}^{-1}{\mathcal{H}_{\mathcal{R}}}^{-1}{\mathcal{H}_{\mathcal{L}}}^{-1}\mathbf{H}\rangle\\ &= \frac{2}{\gamma b\left(\frac{n}{b}-s\right)^2}\langle\mathbf{\Sigma},{\mathcal{T}_b}^{-1}\mathbf{H}^{-1}\rangle\\ &\leq \frac{2}{\gamma b\lambda_{\min}(\mathbf{H})\left(\frac{n}{b}-s\right)^2}\langle\mathbf{I},{\mathcal{T}_b}^{-1}\mathbf{\Sigma}\rangle\\ &\leq \frac{3\,\mathrm{Tr}\,\mathbf{H}^{-1}\mathbf{\Sigma}}{\gamma b\lambda_{\min}(\mathbf{H})\left(\frac{n}{b}-s\right)^2}\end{aligned}$$

which wraps up bounding $\langle\mathbf{H},\mathfrak{T}_4\rangle$. We consider the final term:

$$\begin{aligned}\langle\mathbf{H},\mathfrak{T}_5\rangle &= \|\mathbf{H}\|\,\mathrm{Tr}\left(\mathfrak{T}_5\right)\\ &\leq \sqrt{d}\,\|\mathbf{H}\|\,\|\mathfrak{T}_5\|_F\\ &\leq \sqrt{d}\,\|\mathbf{H}\|\cdot\frac{\rho_b^{\frac{n}{b}-s+1}}{(n-bs)\lambda_{\min}(\mathbf{H})\lambda_{\min}(\mathcal{T}_b)}\|\mathbf{\Sigma}\|_F\\ &\leq \sqrt{d}\,\|\mathbf{H}\|\cdot\frac{2\rho_b^{\frac{n}{b}-s+1}}{(n-bs)\lambda_{\min}(\mathbf{H})^2}\|\mathbf{\Sigma}\|_F\end{aligned}$$

$$\leq 2\sqrt{d}\kappa^2 \cdot \frac{\rho_b^{\frac{n}{b}-s+1}}{(n-bs)} \operatorname{Tr} \mathbf{H}^{-1}\mathbf{\Sigma}$$

which wraps up bounding all terms to establish the proof of Theorem 2. □

### A.3.2 Proof of Theorem 3

*Proof.* We first prove a recursive bound on the error $R(\mathbf{w}_e)$ in epoch $e$ using Theorem 2 for $e \leq \log\left(\frac{n}{bt}\right) - 1$. Since in epoch $e$, we use a minibatch size of $b_e = 2^{e-1}b = 2^{e-1}\frac{R^2 d \left\|(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\right\|_2}{\|\mathbf{H}\|_2 \operatorname{Tr}(\mathcal{H}_\mathcal{L}+\mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}}$, number of initial iterations $s = t-1$ and total samples $n_e \overset{\text{def}}{=} b_e t$, we obtain from Theorem 2 the following bound on the error of the iterate $\mathbf{w}_e$:

$$\begin{aligned}
R(\mathbf{w}_e) - R(\mathbf{w}^*) &\leq \frac{36\sqrt{d}\kappa^2}{\left(\frac{n_e}{b_e}-s\right)^2} \exp\left(-\frac{s}{3\kappa}\right)\{R(\mathbf{w}_{e-1}) - R(\mathbf{w}^*)\} + \exp\left(-\frac{n_e}{3b_e\kappa}\right)\frac{3\kappa^2\sqrt{d}}{\frac{n_e}{b_e}-s}\{R(\mathbf{w}_{e-1}) - R(\mathbf{w}^*)\} \\
&\quad + \frac{3\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n_e - b_e s} + \frac{9\kappa \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{b_e\left(\frac{n_e}{b_e}-s\right)^2} + \frac{2\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{(n_e - b_e s)} \exp\left(-\frac{(\frac{n_e}{b_e}-s+1)}{3\kappa}\right) \\
&\leq 54\sqrt{d}\kappa^2 \exp\left(-\frac{t}{3\kappa}\right)\{R(\mathbf{w}_{e-1}) - R(\mathbf{w}^*)\} + \frac{3\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{e-1}b} \quad \left(\text{ using } b_e = 2^{e-1}b,\ n_e = b_e t,\ s = t-1\right) \\
&\leq \exp\left(-\frac{t}{3\kappa\log(d\kappa)}\right)\{R(\mathbf{w}_{e-1}) - R(\mathbf{w}^*)\} + \frac{3\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{e-1}b} \quad \left(\text{ since } t \geq 16\kappa\log(d\kappa)\right).
\end{aligned}$$

By expanding out the recursion, we obtain:

$$\begin{aligned}
R(\mathbf{w}_e) - R(\mathbf{w}^*) &\leq \exp\left(-\frac{te}{3\kappa\log(d\kappa)}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \sum_{\ell=1}^{e} \frac{3\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{\ell-1}b} \cdot \frac{1}{16^{e-\ell}} \\
&= \exp\left(-\frac{te}{3\kappa\log(d\kappa)}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{3\sigma^2\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{e-1}b} \sum_{\ell=1}^{e} \frac{1}{8^{e-\ell}} \\
&\leq \exp\left(-\frac{te}{3\kappa\log(d\kappa)}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{3\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{e-1}b} \cdot \frac{8}{7} \\
&\leq \exp\left(-\frac{te}{3\kappa\log(d\kappa)}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{16\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{e+1}b}. \qquad (30)
\end{aligned}$$

In the final step of Algorithm 2, we have a minibatch size of $\overline{b} = \frac{n}{2t}$, number of samples $\overline{n} = n/2$ and number of initial iterations $s = \frac{t}{2}$. Plugging this into Theorem 2 again, we obtain the following bound on the final error.

$$\begin{aligned}
R(\overline{\mathbf{w}}) - R(\mathbf{w}^*) &\leq \frac{144\kappa^2\sqrt{d}}{t^2} \exp\left(-\frac{t}{6\kappa}\right)\{R(\mathbf{w}_e) - R(\mathbf{w}^*)\} + \frac{6\kappa^2\sqrt{d}}{t} \exp\left(-\frac{t}{3\kappa}\right)\{R(\mathbf{w}_e) - R(\mathbf{w}^*)\} \\
&\quad + \frac{3\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{\overline{n} - \overline{b}s} + \frac{9\kappa \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{\overline{b}\left(\frac{\overline{n}}{\overline{b}}-s\right)^2} + \frac{2\kappa^2\sqrt{d}\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{(\overline{n} - \overline{b}s)} \exp\left(-\frac{(\frac{\overline{n}}{\overline{b}}-s+1)}{3\kappa}\right) \\
&\leq 3\sqrt{d}\kappa^2 \exp\left(-\frac{t}{3\kappa}\right)\{R(\mathbf{w}_e) - R(\mathbf{w}^*)\} + \frac{4\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}\left(3 + \frac{18\kappa}{t} + 2\kappa^2\sqrt{d}\exp\left(-\frac{t}{6\kappa}\right)\right) \\
&\qquad \left(\text{ using } \overline{b} = \frac{n}{2t},\ \overline{n} = n/2,\ s = t/2\right) \\
&\leq 3\sqrt{d}\kappa^2 \exp\left(-\frac{t}{3\kappa}\right)\{R(\mathbf{w}_e) - R(\mathbf{w}^*)\} + \frac{20\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n} \quad \left(\text{ since } t \geq 16\kappa\log(d\kappa)\right). \qquad (31)
\end{aligned}$$

Since in the final iteration, we have $e = \log \frac{n}{bt} - 1$, using (30), we have

$$R(\mathbf{w}_e) - R(\mathbf{w}^*) \le \exp\left(-\frac{te}{3\kappa \log(d\kappa)}\right)\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{16\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{2^{e+1}b}$$
$$\le \left(\frac{2bt}{n}\right)^{\frac{t}{3\kappa\log(d\kappa)}}\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{16t\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}.$$

Plugging the above inequality into (31), we obtain

$$R(\overline{\mathbf{w}}) - R(\mathbf{w}^*) \le 3\sqrt{d}\kappa^2 \exp\left(-\frac{t}{3\kappa}\right)\{R(\mathbf{w}_e) - R(\mathbf{w}^*)\} + \frac{20 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}$$
$$\le 3\sqrt{d}\kappa^2 \exp\left(-\frac{t}{3\kappa}\right)\left(\left(\frac{2bt}{n}\right)^{\frac{t}{3\kappa\log(d\kappa)}}\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{16t\sqrt{d}\kappa^2 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}\right) + \frac{20 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n}$$
$$\le \left(\frac{2bt}{n}\right)^{\frac{t}{3\kappa\log(d\kappa)}}\{R(\mathbf{w}_0) - R(\mathbf{w}^*)\} + \frac{40 \operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma}}{n} \quad (\text{ since } t > 16\kappa\log(d\kappa))\,.$$

This proves the theorem. □

## A.4 Proof of Lemma 4

For this problem instance, we begin by noting that $(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}$ is diagonal as well, with entries:

$$\{(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\}_{ii} = \frac{1}{2}\{\mathbf{H}^{-1}\mathbf{\Sigma}\}_{ii} = \begin{cases} 1/2 & \text{if } i = 1 \\ 1/2(d-1) & \text{if } i > 1 \end{cases}$$

Let us consider the case with batch size $b = 1$. With the appropriate choice of step size $\gamma$ that ensure contracting operators, we require considering $\operatorname{Tr}[\mathcal{T}_b^{-1}\mathbf{\Sigma}]$ as in equation 29, which corresponds to bounding the leading order term in the variance. We employ the taylor's expansion (just as in claim 1 of lemma 15) to expand the term of interest $\mathcal{T}_b^{-1}\mathbf{\Sigma}$:

$$\mathcal{T}_b^{-1}\mathbf{\Sigma} = \sum_{i=0}^{\infty}\left(\gamma(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathcal{M}\right)^i(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}$$
$$= (\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma} + \sum_{i=1}^{\infty}\left(\gamma(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathcal{M}\right)^i(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}$$
$$\Rightarrow \operatorname{Tr}\mathcal{T}_b^{-1}\mathbf{\Sigma} = \operatorname{Tr}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma} + \sum_{i=1}^{\infty}\operatorname{Tr}\left[\left(\gamma(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathcal{M}\right)^i(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\right]$$
$$\operatorname{Tr}\mathcal{T}_b^{-1}\mathbf{\Sigma} = \frac{1}{2}\operatorname{Tr}\mathbf{H}^{-1}\mathbf{\Sigma} + \sum_{i=1}^{\infty}\operatorname{Tr}\left[\left(\gamma(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathcal{M}\right)^i(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}\right]$$

We observe that the term corresponding to $i = 0$ works out regardless of the choice of stepsize $\gamma$; we then switch our attention to the second term, i.e., the term corresponding to $i = 1$:

$$\operatorname{Tr}\left(\gamma(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathcal{M}\right)(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma} = \frac{d+2}{4}\cdot\operatorname{Tr}\mathbf{\Sigma}$$

We require that this term should be $\le \operatorname{Tr}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}$, implying,

$$\gamma < \frac{4\operatorname{Tr}(\mathcal{H}_\mathcal{L} + \mathcal{H}_\mathcal{R})^{-1}\mathbf{\Sigma}}{(d+2)\operatorname{Tr}\mathbf{\Sigma}}$$

For this example, we observe that this yields $\gamma < \frac{4}{(d+2)(1+\frac{1}{d})}$, which clearly is off by a factor $d$ compared to the realizable case which requires $\gamma < \frac{d}{(d+2)(1+\frac{1}{d})}$, thus indicating a clear separation between the step sizes required by SGD for the realizable and agnostic noise cases.